# Larger is Better: The Effect of Learning Rates Enjoyed by Stochastic Optimization with Progressive Variance Reduction

Fanhua Shang

Department of Computer Science and Engineering, The Chinese University of Hong Kong
fhshang@cse.cuhk.edu.hk

April 18, 2017

arXiv:1704.04966v1 [cs.LG] 17 Apr 2017

**Abstract**

In this paper, we propose a simple variant of the original stochastic *variance reduction* gradient (SVRG) [1], where hereafter we refer to as the *variance reduced stochastic gradient descent* (VR-SGD). Different from the choices of the *snapshot point* and *starting point* in SVRG and its proximal variant, Prox-SVRG [2], the two vectors of each epoch in VR-SGD are set to the *average* and *last iterate* of the previous epoch, respectively. This setting allows us to use much larger learning rates or step sizes than SVRG, e.g., $3/(7L)$ for VR-SGD vs. $1/(10L)$ for SVRG, and also makes our convergence analysis more challenging. In fact, a larger learning rate enjoyed by VR-SGD means that the variance of its stochastic gradient estimator asymptotically approaches zero more rapidly. Unlike common stochastic methods such as SVRG and proximal stochastic methods such as Prox-SVRG, we design two different update rules for *smooth* and *non-smooth* objective functions, respectively. In other words, VR-SGD can tackle non-smooth and/or non-strongly convex problems directly without using any reduction techniques such as quadratic regularizers. Moreover, we analyze the *convergence properties* of VR-SGD for *strongly convex* problems, which show that VR-SGD attains a *linear* convergence rate. We also provide the *convergence guarantees* of VR-SGD for *non-strongly convex* problems. Experimental results show that the performance of VR-SGD is significantly better than its counterparts, SVRG and Prox-SVRG, and it is also much better than the *best known* stochastic method, Katyusha [3].

**Index Terms**

Stochastic optimization, stochastic gradient descent (SGD), proximal stochastic gradient, variance reduction, iterate averaging, snapshot and starting points, strongly convex and non-strongly convex, smooth and non-smooth

---

- *All the codes of VR-SGD and some related variance reduced stochastic methods can be downloaded from the author's website: https://sites.google.com/site/fanhua217/publications.*

## 1 INTRODUCTION

In this paper, we focus on the following *composite* convex optimization problem:

$$\min_{x\in\mathbb{R}^d} F(x) \stackrel{\text{def}}{=} \frac{1}{n}\sum_{i=1}^{n} f_i(x) + g(x) \tag{1}$$

where $f(x) := \frac{1}{n}\sum_{i=1}^{n} f_i(x)$, $f_i(x) : \mathbb{R}^d \to \mathbb{R}$, $i = 1, \ldots, n$ are the smooth convex functions, and $g(x)$ is a relatively simple (but possibly non-differentiable) convex function (referred to as a regularizer). The formulation (1) arises in many places in machine learning, signal processing, data science, statistics and operations research, such as *regularized empirical risk minimization* (ERM). For instance, one popular choice of the component function $f_i(\cdot)$ in binary classification problems is the logistic loss, i.e., $f_i(x) = \log(1 + \exp(-b_i a_i^T x))$, where $\{(a_1, b_1), \ldots, (a_n, b_n)\}$ is a collection of training examples, and $b_i \in \{\pm 1\}$. Some popular choices for the regularizer include the $\ell_2$-norm regularizer (i.e., $g(x) = (\lambda_1/2)\|x\|^2$), the $\ell_1$-norm regularizer (i.e., $g(x) = \lambda_2\|x\|_1$), and the elastic-net regularizer (i.e., $g(x) = (\lambda_1/2)\|x\|^2 + \lambda_2\|x\|_1$), where $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$ are two regularization parameters. So far examples of some other applications include deep neural networks [1], [4], [5], [6], group Lasso [7], [8], sparse learning and coding [9], [10], phase retrieval [11], matrix completion [12], [13], conditional random fields [14], eigenvector computation [15], [16] such as principal component analysis (PCA) and singular value decomposition (SVD), generalized eigen-decomposition and canonical correlation analysis (CCA) [17].

### 1.1 Stochastic Gradient Descent

In this paper, we are especially interested in developing efficient algorithms to solve regularized ERM problems involving a large sum of $n$ component functions. The standard and effective method for solving Problem (1) is the (proximal) gradient descent (GD) method, including accelerated proximal gradient (APG) [18], [19], [20], [21]. For *smooth* objective functions, the update rule of GD is

$$x_{k+1} = x_k - \eta_k \left[\frac{1}{n}\sum_{i=1}^{n} \nabla f_i(x_k) + \nabla g(x_k)\right] \tag{2}$$

for $k = 1, 2, \ldots$, where $\eta_k > 0$ is commonly referred to as the step-size in optimization or the learning rate in machine learning. When the regularizer $g(\cdot)$ is *non-smooth*, e.g., the $\ell_1$-norm regularizer, we need to introduce the following proximal operator into (2),

$$x_{k+1} = \text{Prox}_{\eta_k, g}(y_k) := \arg\min_{x\in\mathbb{R}^d} (1/2\eta_k)\cdot\|x - y_k\|^2 + g(x) \tag{3}$$

where $y_k = x_k - (\eta_k/n)\sum_{i=1}^{n} \nabla f_i(x_k)$. The GD methods mentioned above have been proven to achieve linear convergence for *strongly convex* problems, and APG attains the optimal convergence rate of $\mathcal{O}(1/T^2)$ for *non-strongly convex* problems, where $T$ denotes the number of iterations. However, the per-iteration cost of all the batch (or deterministic) methods is $O(nd)$, which is expensive.

Instead of evaluating the full gradient of $f(\cdot)$ at each iteration, an effective alternative is the stochastic (or incremental) gradient descent (SGD) method [22]. SGD only evaluates the gradient of a single component function at each iteration, thus it has much *lower* per-iteration cost, $O(d)$, and has been successfully applied to many large-scale learning problems [4], [23], [24], [25]. The update rule of SGD is formulated as follows:

$$x_{k+1} = x_k - \eta_k[\nabla f_{i_k}(x_k) + \nabla g(x_k)] \tag{4}$$

where $\eta_k \propto 1/k$, and the index $i_k$ is chosen uniformly at random from $\{1, \ldots, n\}$. Although the expectation of $\nabla f_{i_k}(x_k)$ is an *unbiased* estimation for $\nabla f(x_k)$, i.e., $\mathbb{E}[\nabla f_{i_k}(x_k)] = \nabla f(x_k)$, the variance of the stochastic gradient estimator $\nabla f_{i_k}(x_k)$ may be large due to the variance of random sampling [1]. Thus, stochastic gradient estimators are also called "noisy gradients", and we need to gradually reduce its step size, leading to slow convergence. In particular, even under the strongly convex condition, standard SGD attains a slower *sub-linear* convergence rate of $\mathcal{O}(1/T)$ [26], [27].

### 1.2 Accelerated SGD

Recently, many SGD methods with *variance reduction* techniques were proposed, such as stochastic average gradient (SAG) [28], stochastic variance reduced gradient (SVRG) [1], stochastic dual coordinate ascent (SDCA) [29], SAGA [30], stochastic primal-dual coordinate (SPDC) [31], and their proximal variants, such as Prox-SAG [32], Prox-SVRG [2] and Prox-SDCA [33]. All these accelerated SGD methods can use a constant step size $\eta$ instead of diminishing step sizes for SGD, and fall into the following three categories: *primal* methods such as SVRG and SAGA, *dual* methods such as SDCA, and *primal-dual* methods such as SPDC. In essence, many of primal methods use the full gradient at the snapshot $\widetilde{x}$ or average gradients to progressively reduce the variance of stochastic gradient estimators, as well as dual and primal-dual methods, *which leads to a revolution in the area of first-order methods* [34]. Thus, they are also known as the *hybrid gradient descent* method [35] or *semi-stochastic gradient descent* method [36]. In particular, under the strongly convex condition, the accelerated SGD methods enjoy linear convergence rates and the overall complexity of $\mathcal{O}((n+L/\mu)\log(1/\epsilon))$ to obtain an $\epsilon$-suboptimal solution, where each $f_i(\cdot)$ is $L$-smooth and $g(\cdot)$ is $\mu$-strongly convex. The complexity bound shows that they converge significantly faster than deterministic APG methods, whose complexity is $\mathcal{O}((n\sqrt{L/\mu})\log(1/\epsilon))$ [36].

SVRG [1] and its proximal variant, Prox-SVRG [2], are particularly attractive because of their *low storage* requirement compared with other stochastic methods such as SAG, SAGA and SDCA, which require to store all the gradients of the $n$ component functions $f_i(\cdot)$ or dual variables. At the beginning of each epoch in SVRG, the full gradient $\nabla f(\widetilde{x})$ is computed at the *snapshot point* $\widetilde{x}$, which is updated periodically. The update rule for the *smooth* optimization problem (1) is given by

$$\widetilde{\nabla} f_{i_k}(x_k) = \nabla f_{i_k}(x_k) - \nabla f_{i_k}(\widetilde{x}) + \nabla f(\widetilde{x}), \tag{5a}$$

$$x_{k+1} = x_k - \eta[\widetilde{\nabla} f_{i_k}(x_k) + \nabla g(x_k)]. \tag{5b}$$

When $g(\cdot) \equiv 0$, the update rule in (5b) becomes the original one in [1], i.e., $x_{k+1} = x_k - \eta\widetilde{\nabla} f_{i_k}(x_k)$. It is not hard to verify that the variance of the SVRG estimator $\widetilde{\nabla} f_{i_k}(x_k)$, i.e., $\mathbb{E}\|\widetilde{\nabla} f_{i_k}(x_k) - \nabla f(x_k)\|^2$, can be much smaller than that of the SGD estimator $\nabla f_{i_k}(x_k)$, i.e., $\mathbb{E}\|\nabla f_{i_k}(x_k) - \nabla f(x_k)\|^2$. However, for *non-strongly convex* problems, the accelerated SGD methods mentioned above converge much slower than batch APG methods such as FISTA [21], namely, $\mathcal{O}(1/T)$ vs. $\mathcal{O}(1/T^2)$.

More recently, many *acceleration* techniques were proposed to further speed up those variance-reduced stochastic methods mentioned above. These techniques mainly include the Nesterov's acceleration technique used in [24], [37], [38], [39], [40], reducing the number of gradient calculations in the early iterations [34], [41], [42], the projection-free property of the conditional gradient method (also known as the Frank-Wolfe algorithm [43]) as in [44], the stochastic sufficient decrease technique [45], and the momentum acceleration trick in [3], [34], [46]. More specifically, [40] proposed an accelerating Catalyst framework and achieved the complexity of $\mathcal{O}((n+\sqrt{nL/\mu})\log(L/\mu)\log(1/\epsilon))$ for strongly convex problems. [3] and [46] proved that their accelerated methods can attain the best known complexity of $\mathcal{O}(n\log(1/\epsilon)+\sqrt{nL/\epsilon})$ for non-strongly convex problems. The overall complexity matches the theoretical upper bound provided in [47]. Katyusha [3] and point-SAGA [48] achieve the best-known complexity of $\mathcal{O}((n+\sqrt{nL/\mu})\log(1/\epsilon))$ for strongly convex problems, which is identical to the upper complexity bound in [47]. That is, *Katyusha* is the *best known* stochastic optimization method for both *strongly convex* and *non-strongly convex* problems. Its proximal gradient update rules are formulated as follows:

$$x_{k+1} = w_1 y_k + w_2 \widetilde{x} + (1 - w_1 - w_2) z_k, \tag{6a}$$

$$y_{k+1} = \arg\min_{y\in\mathbb{R}^d} \left\{ \frac{1}{2\eta}\|y - y_k\|^2 + y^T \widetilde{\nabla} f_{i_k}(x_{k+1}) + g(y) \right\}, \tag{6b}$$

$$z_{k+1} = \arg\min_{z\in\mathbb{R}^d} \left\{ \frac{3L}{2}\|z - x_{k+1}\|^2 + z^T \widetilde{\nabla} f_{i_k}(x_{k+1}) + g(z) \right\} \tag{6c}$$

where $w_1, w_2 \in [0, 1]$ are two momentum parameters. To eliminate the need for parameter tuning, $\eta$ is set to $1/(3w_1 L)$, and $w_2$ is fixed to $0.5$ in [3]. Unfortunately, most of the accelerated methods mentioned above, including Katyusha, require at least two auxiliary variables and two momentum parameters, which lead to complicated algorithm design and high per-iteration complexity [34].

### 1.3 Our Contributions

From the above discussion, one can see that most of accelerated stochastic variance reduction methods such as [3], [34], [37], [42], [44], [45] and applications such as [9], [10], [13], [15], [16], [49] are based on the *stochastic variance reduced gradient* (SVRG) method [1]. Thus, any key improvement on SVRG is very important for the research of stochastic optimization. In this paper, we propose a simple variant of the original SVRG [1], which is referred to as the *variance reduced stochastic gradient descent* (VR-SGD). The *snapshot point* and *starting point* of each epoch in VR-SGD are set to the *average* and *last iterate* of the previous epoch, respectively. Different from the settings of SVRG and Prox-SVRG [2] (i.e., the last iterate for the two points of the former, while the average of the previous epoch for those of the latter), the two points in VR-SGD are different, which makes our convergence analysis *more challenging* than SVRG and Prox-SVRG. Our empirical results show that the performance of VR-SGD is significantly better than its counterparts, SVRG and Prox-SVRG. Impressively, VR-SGD with a *sufficiently large* learning rate performs much better than the *best known* stochastic method, Katyusha [3]. The main contributions of this paper are summarized below.

- The *snapshot point* and *starting point* of VR-SGD are set to two different vectors. That is, for all epochs, except the first one, $\widetilde{x}^s = \frac{1}{m-1}\sum_{k=1}^{m-1} x_k^s$ (denoted by *Option I*) or $\widetilde{x}^s = \frac{1}{m}\sum_{k=1}^{m} x_k^s$ (denoted by *Option II*), and $x_0^{s+1} = x_m^s$. In particular, we find that the setting of VR-SGD allows us take much *larger learning rates* or step sizes than SVRG, e.g., $3/(7L)$ vs. $1/(10L)$, and thus significantly speeds up the convergence of SVRG and Prox-SVRG in practice. Moreover, VR-SGD has an advantage over SVRG in terms of *robustness of learning rate selection*.
- Different from *proximal stochastic gradient* methods, e.g., Prox-SVRG and Katyusha, which have a unified update rule for the two cases of *smooth* and *non-smooth* objectives (see Section 2.2 for details), VR-SGD employs *two different update rules* for the two cases, respectively, as in (12) and (13) below. Empirical results show that gradient update rules as in (12) for *smooth* optimization problems are better choices than proximal update formulas as in (10).
- Finally, we theoretically analyze the *convergence properties* of VR-SGD with Option I or Option II for *strongly convex* problems, which show that VR-SGD attains a *linear* convergence rate. We also give the *convergence guarantees* of VR-SGD with Option I or Option II for *non-strongly convex* objective functions.

## 2 Preliminary and Related Work

Throughout this paper, we use $\|\cdot\|$ to denote the $\ell_2$-norm (also known as the standard Euclidean norm), and $\|\cdot\|_1$ is the $\ell_1$-norm, i.e., $\|x\|_1 = \sum_{i=1}^d |x_i|$. $\nabla f(\cdot)$ denotes the full gradient of $f(\cdot)$ if it is differentiable, or $\partial f(\cdot)$ the subgradient if $f(\cdot)$ is only Lipschitz continuous. For each epoch $s \in [S]$ and inner iteration $k \in \{0, 1, \ldots, m-1\}$, $i_k^s \in [n]$ is the random chosen index. We mostly focus on the case of Problem (1) when each component function $f_i(\cdot)$ is $L$-smooth[1], and $F(\cdot)$ is $\mu$-strongly convex. The two common assumptions are defined as follows.

### 2.1 Basic Assumptions

**Assumption 1** (Smoothness). *Each convex function $f_i(\cdot)$ is $L$-smooth, that is, there exists a constant $L > 0$ such that for all $x, y \in \mathbb{R}^d$,*

$$\|\nabla f_i(x) - \nabla f_i(y)\| \leq L\|x - y\|. \tag{7}$$

**Assumption 2** (Strong Convexity). *The convex function $F(x)$ is $\mu$-strongly convex, i.e., there exists a constant $\mu > 0$ such that for all $x, y \in \mathbb{R}^d$,*

$$F(y) \geq F(x) + \langle \nabla F(x),\, y - x\rangle + \frac{\mu}{2}\|x - y\|^2. \tag{8}$$

Note that when the regularizer $g(\cdot)$ is *non-smooth*, the inequality in (8) needs to be revised by simply replacing the gradient $\nabla F(x)$ in (8) with an arbitrary sub-gradient of $F(\cdot)$ at $x$. In contrast, for a non-strongly convex or *general convex* function, the inequality in (8) can always be satisfied with $\mu = 0$.

1. Actually, we can extend all the theoretical results in this paper for the case, when the gradients of all component functions have the same Lipschitz constant $L$, to the more general case, when some $f_i(\cdot)$ have different degrees of smoothness.

**Algorithm 1** SVRG (Option I) and Prox-SVRG (Option II)

**Input:** The number of epochs $S$, the number of iterations $m$ per epoch, and step size $\eta$.

**Initialize:** $\widetilde{x}^0$.

1: **for** $s = 1, 2, \ldots, S$ **do**
2: $\quad x_0^s = \widetilde{x}^{s-1}$; $\quad$ % *Initiate the variable* $x_0^s$
3: $\quad \widetilde{\mu}^s = \frac{1}{n}\sum_{i=1}^{n} \nabla f_i(\widetilde{x}^{s-1})$; $\quad$ % *Compute the full gradient*
4: $\quad$ **for** $k = 0, 1, \ldots, m-1$ **do**
5: $\quad\quad$ Pick $i_k^s$ uniformly at random from $[n]$;
6: $\quad\quad \widetilde{\nabla} f_{i_k^s}(x_k^s) = \nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1}) + \widetilde{\mu}^s$; $\quad$ % *The stochastic gradient estimator*
7: $\quad\quad$ Option I: $x_{k+1}^s = x_k^s - \eta\left[\widetilde{\nabla} f_{i_k^s}(x_k^s) + \nabla g(x_k^s)\right]$, $\quad$ % *Smooth case of* $g(\cdot)$
$\quad\quad\quad$ or $x_{k+1}^s = \mathrm{Prox}_{\eta, g}\left(x_k^s - \eta \widetilde{\nabla} f_{i_k^s}(x_k^s)\right)$; $\quad$ % *Non-smooth case of* $g(\cdot)$
8: $\quad\quad$ Option II: $x_{k+1}^s = \arg\min_{y \in \mathbb{R}^d}\left\{g(y) + y^T \widetilde{\nabla} f_{i_k^s}(x_k^s) + \frac{1}{2\eta}\|y - x_k^s\|^2\right\}$; $\quad$ % *Proximal update*
9: $\quad$ **end for**
10: $\quad$ Option I: $\widetilde{x}^s = x_m^s$; $\quad$ % *Last iterate for snapshot* $\widetilde{x}$
11: $\quad$ Option II: $\widetilde{x}^s = \frac{1}{m}\sum_{k=1}^{m} x_k^s$; $\quad$ % *Iterate averaging for snapshot* $\widetilde{x}$
12: **end for**

**Output:** $\widetilde{x}^S$

### 2.2 Related Work

To speed up standard and proximal SGD methods, many variance reduced stochastic methods [28], [29], [30], [35] have been proposed for some special cases of Problem (1). In the case when each $f_i(x)$ is $L$-smooth, $f(x)$ is $\mu$-strongly convex, and $g(x) \equiv 0$, Roux *et al.* [28] proposed a stochastic average gradient (SAG) method, which attains a linear convergence rate. However, SAG needs to store all gradients as well as other incremental aggregated gradient methods such as SAGA [30], so that $O(nd)$ storage is required in general problems [41]. Similarly, SDCA [29] requires storage of all dual variables [1], which scales as $O(n)$. In contrast, SVRG [1], as well as its proximal variant, Prox-SVRG [2], has the similar convergence rate to SAG and SDCA but without the memory requirements of all gradients and dual variables. In particular, the SVRG estimator in (5a) (independently introduced in [1], [35]) may be the most popular choice for *stochastic gradient estimators*. Besides, other stochastic gradient estimators include the SAGA estimator in [30] and the stochastic recursive gradient estimator in [50]. Although the original SVRG in [1] only has convergence guarantees for a special case of Problem (1), when each $f_i(x)$ is $L$-smooth, $f(x)$ is $\mu$-strongly convex, and $g(x) \equiv 0$, one can extend SVRG to the proximal setting by introducing the proximal operator in (3), as shown in Line 7 of Algorithm 1. In other words, when $g(\cdot)$ is *non-smooth*, the update rule of SVRG becomes

$$x_{k+1}^s = \arg\min_{x \in \mathbb{R}^d}\left\{(1/2\eta)\cdot\|x - [x_k^s - \eta\widetilde{\nabla} f_{i_k^s}(x_k^s)]\|^2 + g(x)\right\}. \tag{9}$$

Some researchers [8], [51], [52] have borrowed some *variance reduction* techniques into ADMM for minimizing convex composite objective functions subject to an equality constraint. [15], [16], [17] applied efficient stochastic solvers to compute leading eigenvectors of a symmetric matrix or generalized eigenvectors of two symmetric matrices. The first such method is VR-PCA by Shamir [15], and the convergence properties of the VR-PCA algorithm for such a non-convex problem are also provided. Garber *et al.* [16] analyzed the convergence rate of SVRG when $f(\cdot)$ is a convex function that is a sum of non-convex component functions. Moreover, [6] and [53] proved that SVRG and SAGA with minor modifications can converge asymptotically to a stationary point of non-convex finite-sum problems. Some distributed variants [54], [55] of accelerated SGD methods have also been proposed.

An important class of stochastic methods is the *proximal stochastic gradient* (Prox-SG) method, such as Prox-SVRG [2], SAGA [30], and Katyusha [3]. Different from standard variance reduction SGD methods such as SVRG, which have a

stochastic gradient update as in (5b), the Prox-SG method has a unified update rule for both smooth and non-smooth cases of $g(\cdot)$. For instance, the update rule of Prox-SVRG [2] is formulated as follows:

$$x_{k+1}^s = \arg\min_{y\in\mathbb{R}^d} \left\{ g(y) + y^T \widetilde{\nabla} f_{i_k^s}(x_k^s) + \frac{1}{2\eta}\|y - x_k^s\|^2 \right\}. \tag{10}$$

For the sake of completeness, the details of Prox-SVRG [2] are shown in Algorithm 1 with Option II. When $g(\cdot)$ is the widely used $\ell_2$-norm regularizer, i.e., $g(\cdot) = (\lambda_1/2)\|\cdot\|^2$, the proximal update formula in (10) becomes

$$x_{k+1}^s = \frac{1}{1+\lambda_1\eta}\left[x_k^s - \eta\widetilde{\nabla} f_{i_k^s}(x_k^s)\right]. \tag{11}$$

## 3 Variance-Reduced Stochastic Gradient Descent

In this section, we propose an efficient variance reduced stochastic gradient descent (VR-SGD) method with *iterate averaging*. Different from the choices of the snapshot and starting points in SVRG [1] and Prox-SVRG [2], the two vectors of each epoch in VR-SGD are set to the average and last iterate of the previous epoch, respectively. Unlike common stochastic gradient methods such as SVRG and proximal stochastic gradient methods such as Prox-SVRG, we design two different update rules for smooth and non-smooth objective functions, respectively.

### 3.1 Iterate Averaging

Like SVRG and Katyusha, VR-SGD is also divided into $S$ epochs, and each epoch consists of $m$ stochastic gradient steps, where $m$ is usually chosen to be $\Theta(n)$ as in [1], [2], [3]. Within each epoch, we need to compute the full gradient $\nabla f(\widetilde{x}^s)$ at the snapshot $\widetilde{x}^s$ and use it to define the variance reduced stochastic gradient estimator $\widetilde{\nabla} f_{i_k}(x_k^s)$ as in [1]. Unlike SVRG whose snapshot point is set to the last iterate of the previous epoch, the *snapshot* $\widetilde{x}^s$ of each epoch in VR-SGD is set to the *average* of the previous epoch, e.g., $\widetilde{x}^s = \frac{1}{m-1}\sum_{k=1}^{m-1} x_k^s$ in Option I of Algorithm 2, which leads to better robustness to gradient noise[2], as also suggested in [34], [45], [59]. In fact, the choice of Option II in Algorithm 2, i.e., $\widetilde{x}^s = \frac{1}{m}\sum_{k=1}^{m} x_k^s$, also works well in practice, as shown in Fig. 1. Therefore, we provide the convergence guarantees for both our algorithms (including Algorithm 2) with Option I and our algorithms with Option II in the next section. In particular, we find that the one of the effects of the choice in Option I or Option II of Algorithm 2 is to allow taking much larger learning rates or step sizes than SVRG in practice, e.g., $3/(7L)$ for VR-SGD vs. $1/(10L)$ for SVRG (see Fig. 2 and Section 5.3 for details). This is the main reason why VR-SGD converges significantly faster than SVRG. Actually, a larger learning rate enjoyed by VR-SGD means that the variance of its stochastic gradient estimator goes asymptotically to zero faster.

Unlike Prox-SVRG [2] whose starting point is initialized to the average of the previous epoch, the starting point $x_0^{s+1}$ of each epoch in VR-SGD is set to the last iterate $x_m^s$ of the previous epoch. That is, the last iterate of the previous epoch becomes the new starting point in VR-SGD, while those of Prox-SVRG are completely different, thereby leading to relatively slow convergence in general. It is clear that both the starting point and snapshot point of each epoch in the original SVRG [1] are set to the last iterate of the previous epoch[3], while the two points of Prox-SVRG [2] are set to the average of the previous epoch (also suggested in [1]). Different from the settings in SVRG and Prox-SVRG, the starting and snapshot points in VR-SGD are set to the two different vectors mentioned above, which makes the convergence analysis of VR-SGD more challenging than SVRG and Prox-SVRG, as shown in Section 4.

2. It should be emphasized that the noise introduced by random sampling is inevitable, and generally slows down the convergence speed in a sense. However, SGD and its variants are probably the most used optimization algorithms for deep learning [56]. In particular, [57] has shown that by adding noise at each step, noisy gradient descent can escape the saddle points efficiently and converge to a local minimum of non-convex optimization problems, as well as the application of deep neural networks in [58].

3. Note that the theoretical convergence of the original SVRG [1] relies on its Option II, i.e., both $\widetilde{x}^s$ and $x_0^{s+1}$ are set to $x_k^s$, where $k$ is randomly chosen from $\{1, 2, \ldots, m\}$. However, the empirical results in [1] suggest that Option I is a better choice than its Option II, and the convergence guarantee of SVRG with Option I for strongly convex objective functions is provided in [60].

**Algorithm 2** VR-SGD for strongly convex objectives

**Input:** The number of epochs $S$, the number of iterations $m$ per epoch, and step size $\eta$.

**Initialize:** $x_0^1 = \widetilde{x}^0$.

**for** $s = 1, 2, \ldots, S$ **do**
  $\widetilde{\mu}^s = \frac{1}{n}\sum_{i=1}^{n} \nabla f_i(\widetilde{x}^{s-1})$; % *Compute the full gradient*
  **for** $k = 0, 1, \ldots, m-1$ **do**
    Pick $i_k^s$ uniformly at random from $[n]$;
    $\widetilde{\nabla} f_{i_k^s}(x_k^s) = \nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1}) + \widetilde{\mu}^s$; % *The stochastic gradient estimator*
    $x_{k+1}^s = x_k^s - \eta\Big[\widetilde{\nabla} f_{i_k^s}(x_k^s) + \nabla g(x_k^s)\Big]$, % *Smooth case of* $g(\cdot)$
    or $x_{k+1}^s = \mathrm{Prox}_{\eta,\, g}\Big(x_k^s - \eta\widetilde{\nabla} f_{i_k^s}(x_k^s)\Big)$; % *Non-smooth case of* $g(\cdot)$
  **end for**
  Option I: $\widetilde{x}^s = \frac{1}{m-1}\sum_{k=1}^{m-1} x_k^s$; % *Iterate averaging for snapshot* $\widetilde{x}$
  Option II: $\widetilde{x}^s = \frac{1}{m}\sum_{k=1}^{m} x_k^s$; % *Iterate averaging for snapshot* $\widetilde{x}$
  $x_0^{s+1} = x_m^s$; % *Initiate* $x_0^{s+1}$ *for the next epoch*
**end for**

**Output:** $\widehat{x}^S = \widetilde{x}^S$ if $F(\widetilde{x}^S) \leq F(\frac{1}{S}\sum_{s=1}^{S}\widetilde{x}^s)$, and $\widehat{x}^S = \frac{1}{S}\sum_{s=1}^{S}\widetilde{x}^s$ otherwise.

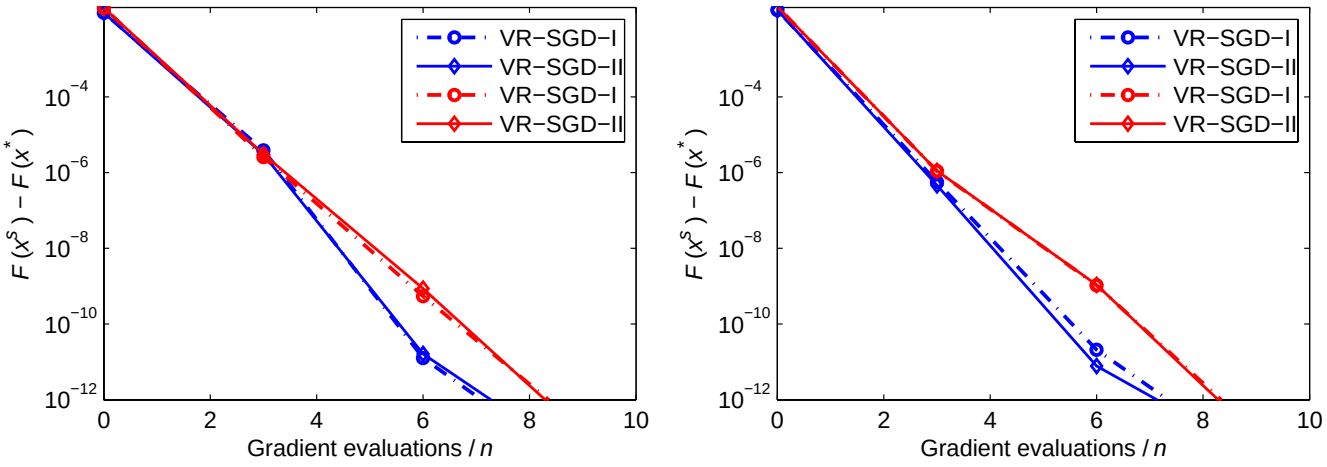


(a) Logistic regression (b) Ridge regression

Fig. 1. Comparison of VR-SGD with Option I (denoted by VR-SGD-I) and Option II (denoted by VR-SGD-II) for solving $\ell_2$-norm (i.e., $(\lambda/2)\|\cdot\|^2$) regularized logistic regression and ridge regression problems on the Covtype data set. In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes. Note that the blue lines stand for the results where $\lambda = 10^{-4}$, while the red lines correspond to the results where $\lambda = 10^{-5}$ (best viewed in colors).

### 3.2 The Algorithm for Strongly Convex Objectives

In this part, we propose an efficient VR-SGD algorithm to solve *strongly convex* objective functions, as outlined in **Algorithm 2**. It is well known that the original SVRG [1] only works for the case of smooth and strongly convex objective functions. However, in many machine learning applications, e.g., elastic net regularized logistic regression, the strongly convex objective function $F(x)$ is non-smooth. To solve this class of problems, the proximal variant of SVRG, Prox-SVRG [2], was subsequently proposed. Unlike SVRG and Prox-SVRG, VR-SGD can not only solve *smooth* objective functions, but directly tackle *non-smooth* ones. That is, when the regularizer $g(x)$ is smooth, e.g., the $\ell_2$-norm regularizer, the update rule of VR-SGD is

$$x_{k+1}^s = x_k^s - \eta[\widetilde{\nabla} f_{i_k^s}(x_k^s) + \nabla g(x_k^s)]. \tag{12}$$

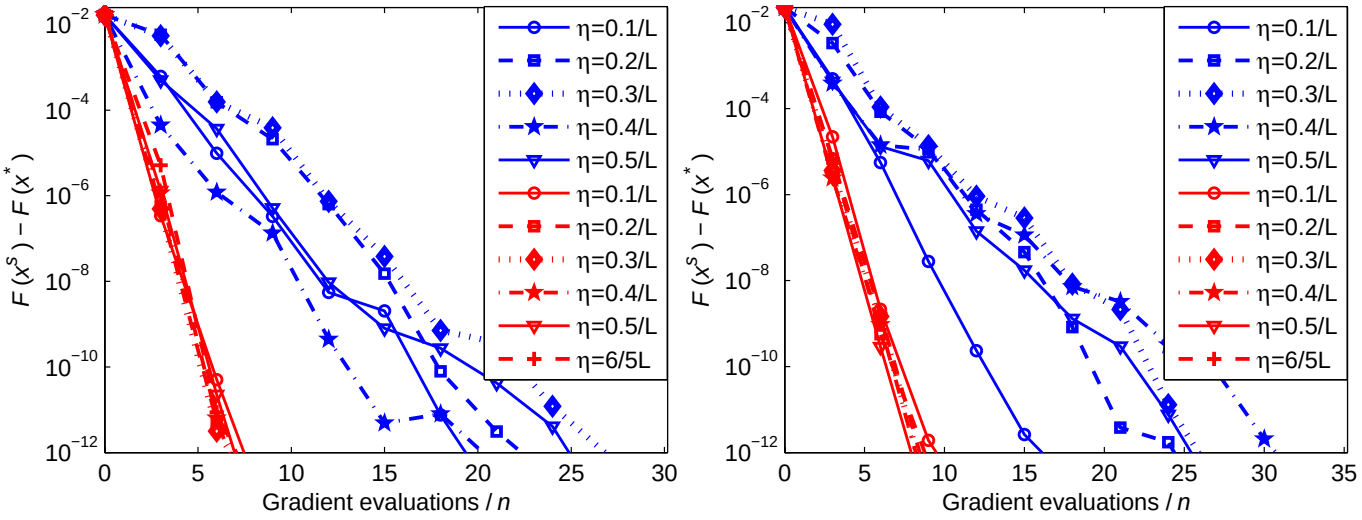


(a) Logistic regression: $\lambda = 10^{-4}$ (left) and $\lambda = 10^{-5}$ (right)

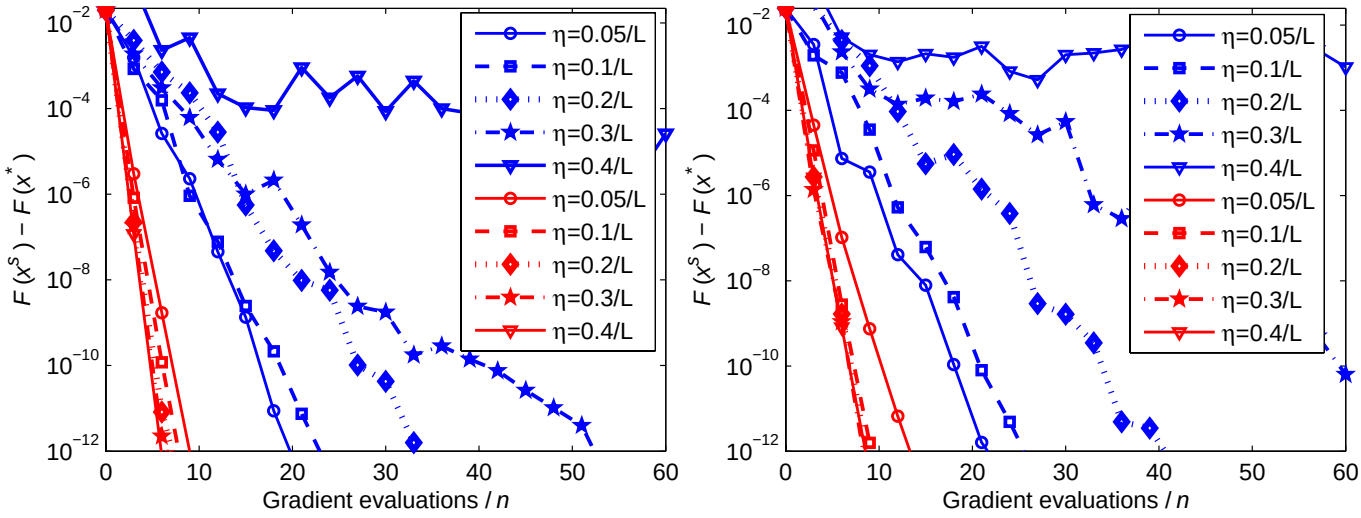


(b) Ridge regression: $\lambda = 10^{-4}$ (left) and $\lambda = 10^{-5}$ (right)

Fig. 2. Comparison of SVRG [1] and VR-SGD with different learning rates for solving $\ell_2$-norm (i.e., $(\lambda/2)\|\cdot\|^2$) regularized logistic regression and ridge regression problems on the Covtype data set. In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes. Note that the *blue* lines stand for the results of SVRG with different learning rates, while the *red* lines correspond to the results of VR-SGD with different learning rates (best viewed in colors).

When $g(x)$ is non-smooth, e.g., the $\ell_1$-norm regularizer, the update rule of VR-SGD becomes

$$x_{k+1}^s = \mathrm{Prox}_{\eta,\, g}\Big(x_k^s - \eta\widetilde{\nabla} f_{i_k^s}(x_k^s)\Big)\,. \tag{13}$$

Different from the proximal stochastic gradient methods such as Prox-SVRG [2], all of which have a unified update rule as in (10) for both the smooth and non-smooth cases of $g(\cdot)$, VR-SGD has two different update rules for the two cases, as stated in (12) and (13). This leads to the following advantage over the Prox-SG methods: the stochastic gradient update rule in (12) usually outperforms the proximal stochastic gradient update rule in (11), as well as the two classes of update rules for Katyusha [3] (see Section 5.4 for details).

Fig. 2 demonstrates that VR-SGD has a significant advantage over SVRG in terms of robustness of learning rate selection. That is, VR-SGD yields good performance within the range of the learning rate between $0.1/L$ and $0.4/L$, whereas the performance of SVRG is very *sensitive* to the selection of learning rates. Thus, VR-SGD is convenient to apply in various real-world problems of large-scale machine learning. In fact, VR-SGD can use much larger learning rates than SVRG for logistic regression problems in practice, e.g., $6/(5L)$ for VR-SGD vs. $1/(10L)$ for SVRG, as shown in Fig. 2(a).

### 3.3 The Algorithm for Non-Strongly Convex Objectives

Although many variance reduced stochastic methods have been proposed, most of them, including SVRG and Prox-SVRG, only have convergence guarantees for the case of Problem (1) when the objective function $F(x)$ is strongly convex. However, $F(x)$ may be non-strongly convex in many machine learning applications such as Lasso and $\ell_1$-norm regularized logistic regression. As suggested in [3], [61], this class of problems can be transformed into strongly convex ones by adding a proximal term $(\tau/2)\|x - x_0^s\|^2$, which can be efficiently solved by Algorithm 2. However, the reduction technique may degrade the performance of the involved algorithms both in theory and in practice [42]. Thus, we present an efficient VR-SGD algorithm for directly solving the *non-strongly convex* problem (1), as outlined in **Algorithm** 3.

The main difference between Algorithm 2 and Algorithm 3 is the setting for the learning rate. Similar to Algorithm 2, the learning rate $\eta_s$ of Algorithm 3 can also be fixed to a constant. Inspired by existing accelerated stochastic algorithms [3], [34], the learning rate $\eta_s$ in Algorithm 3 can be gradually *increased*, which in principle leads to faster convergence (see Section 5.5 for details). Different from existing stochastic methods such as Katyusha [3], the update rule of the learning rate $\eta_s$ for **Algorithm** 3 is defined as follows: $\eta_0 = c$, where $c > 0$ is an initial learning rate, and for any $s \geq 1$,

$$\eta_s = \eta_0 / \max\{\alpha,\, 2/(s+1)\} \tag{14}$$

where $0 < \alpha \leq 1$ is a given constant, e.g., $\alpha = 0.2$.

### 3.4 Complexity Analysis

From Algorithms 2 and 3, we can see that the per-iteration cost of VR-SGD is dominated by the computation of $\nabla f_{i_k^s}(x_k^s)$, $\nabla f_{i_k^s}(\widetilde{x}^{s-1})$, and $\nabla g(x_k^s)$ or the proximal update in (13). Thus, the complexity is $O(d)$, as low as that of SVRG [1] and Prox-SVRG [2]. In fact, for some ERM problems, we can save the intermediate gradients $\nabla f_i(\widetilde{x}^{s-1})$ in the computation of $\widetilde{\mu}^s$, which generally requires $O(n)$ additional storage. As a result, each epoch only requires $(n + m)$ component gradient evaluations. In addition, for extremely sparse data, we can introduce the lazy update tricks in [36], [62], [63] to our algorithms, and perform the update steps in (12) and (13) only for the non-zero dimensions of each example, rather than all dimensions. In other words, the per-iteration complexity of VR-SGD can be improved from $O(d)$ to $O(d')$, where $d' \leq d$ is the sparsity of feature vectors. Moreover, VR-SGD has a much lower per-iteration complexity than existing accelerated stochastic variance reduction methods such as Katyusha [3], which have at least two more update rules for additional variables, as shown in (6a)-(6c).

### 3.5 Extensions of VR-SGD

It has been shown in [36], [37] that *mini-batching* can effectively decrease the variance of stochastic gradient estimates. In this part, we first extend the proposed VR-SGD method to the mini-batch setting, as well as its convergence results below. Here, we denote by $b$ the mini-batch size and $I_k^s$ the selected random index set $I_k \subset [n]$ for each outer-iteration $s \in [S]$ and inner-iteration $k \in \{0, 1, \ldots, m-1\}$. The variance reduced stochastic gradient estimator in (5a) becomes

$$\widetilde{\nabla} f_{I_k^s}(x_k^s) = \frac{1}{b} \sum_{i \in I_k^s} \left[\nabla f_i(x_k^s) - \nabla f_i(\widetilde{x}^{s-1})\right] + \nabla f(\widetilde{x}^{s-1})$$

where $I_k^s \subset [n]$ is a mini-batch of size $b$. If some component functions are non-smooth, we can use the proximal operator oracle [61] or the Nesterov's smoothing [64] and homotopy smoothing [65] techniques to smoothen them, and thereby obtain the smoothed approximations of the functions $f_i(x)$. In addition, we can directly extend our algorithms to the non-smooth setting as in [34], e.g., Algorithm 3 in [34].

Considering that each component function $f_i(x)$ maybe have different degrees of smoothness, picking the random index $i_k^s$ from a non-uniform distribution is a much better choice than the commonly used uniform random sampling [66], [67], as well as without-replacement sampling vs. with-replacement sampling [27]. This can be done using the same techniques in [2], [3], i.e., the sampling probabilities for all $f_i(x)$ are proportional to their Lipschitz constants, i.e., $p_i = L_i / \sum_{j=1}^n L_j$.

**Algorithm 3** VR-SGD for non-strongly convex objectives

**Input:** The number of epochs $S$, and the number of iterations $m$ per epoch.

**Initialize:** $x_0^1 = \widetilde{x}^0$, $\alpha > 0$ and $\eta_0$.

**for** $s = 1, 2, \ldots, S$ **do**
  $\widetilde{\mu}^s = \frac{1}{n}\sum_{i=1}^{n} \nabla f_i(\widetilde{x}^{s-1})$; % *Compute full gradient*
  $\eta_s = \eta_0 / \max\{\alpha, 2/(s+1)\}$; % *Compute step sizes*
  **for** $k = 0, 1, \ldots, m-1$ **do**
    Pick $i_k^s$ uniformly at random from $[n]$;
    $\widetilde{\nabla} f_{i_k^s}(x_k^s) = \nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1}) + \widetilde{\mu}^s$; % *The stochastic gradient estimator*
    $x_{k+1}^s = x_k^s - \eta_s \Big[\widetilde{\nabla} f_{i_k^s}(x_k^s) + \nabla g(x_k^s)\Big]$, % *Smooth case of* $g(\cdot)$
    or $x_{k+1}^s = \mathrm{Prox}_{\eta_s, g}\Big(x_k^s - \eta_s \widetilde{\nabla} f_{i_k^s}(x_k^s)\Big)$; % *Non-smooth case of* $g(\cdot)$
  **end for**
  Option I: $\widetilde{x}^s = \frac{1}{m-1}\sum_{k=1}^{m-1} x_k^s$; % *Iterate averaging for snapshot* $\widetilde{x}$
  Option II: $\widetilde{x}^s = \frac{1}{m}\sum_{k=1}^{m} x_k^s$; % *Iterate averaging for snapshot* $\widetilde{x}$
  $x_0^{s+1} = x_m^s$; % *Initiate* $x_0^{s+1}$ *for the next epoch*
**end for**

**Output:** $\widehat{x}^S = \widetilde{x}^S$ if $F(\widetilde{x}^S) \leq F(\frac{1}{S}\sum_{s=1}^{S} \widetilde{x}^s)$, and $\widehat{x}^S = \frac{1}{S}\sum_{s=1}^{S} \widetilde{x}^s$ otherwise.

Moreover, our VR-SGD method can also be combined with other *accelerated* techniques proposed for SVRG. For instance, the epoch length of VR-SGD can be automatically determined by the techniques in [42], [68] instead of a fixed epoch length. We can reduce the number of gradient calculations in the early iterations as in [34], [41], [42], which leads to faster convergence in general. Moreover, we can also introduce the Nesterov's acceleration technique as in [24], [37], [38], [39], [40] and momentum acceleration trick as in [3], [34] to further improve the performance of VR-SGD.

## 4 Convergence Analysis

In this section, we provide the convergence guarantees of VR-SGD for solving both smooth and non-smooth general convex problems. We also extend these results to the mini-batch setting. Moreover, we analyze the convergence properties of VR-SGD for solving both smooth and non-smooth strongly convex objective functions. We first introduce the following lemma which is useful in our analysis.

**Lemma 1** (3-point property, [69]). *Let $\hat{z}$ be the optimal solution of the following problem,*

$$\min_{z \in \mathbb{R}^d} \frac{\tau}{2}\|z - z_0\|^2 + r(z)$$

*where $r(z)$ is a convex function (but possibly non-differentiable), and $\tau \geq 0$. Then for any $z \in \mathbb{R}^d$, the following inequality holds*

$$r(\hat{z}) + \frac{\tau}{2}\|\hat{z} - z_0\|^2 \leq r(z) + \frac{\tau}{2}\|z - z_0\|^2 - \frac{\tau}{2}\|z - \hat{z}\|^2.$$

### 4.1 Convergence Properties for Non-strongly Convex Problems

In this part, we analyze the convergence properties of VR-SGD for solving the more general non-strongly convex problems. Considering that the two proposed algorithms (i.e., Algorithms 2 and 3) have two different update rules for the smooth and non-smooth problems, we give the convergence guarantees of VR-SGD for the two cases as follows.

#### 4.1.1 Smooth Objectives

We first give the convergence guarantee of Problem (1) when the objective function $F(x)$ is *smooth* and non-strongly convex. In order to simplify analysis, we denote $F(x)$ by $f(x)$, that is, $f_i(x) := f_i(x) + g(x)$ for all $i = 1, 2, \ldots, n$.

**Lemma 2** (Variance bound of smooth objectives). *Let $x^*$ be the optimal solution of Problem (1). Suppose Assumption 1 holds, and let $\widetilde{\nabla} f_{i_k^s}(x_k^s) := \nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1}) + \nabla f(\widetilde{x}^{s-1})$. Then the following inequality holds*

$$\mathbb{E}\left[\|\widetilde{\nabla} f_{i_k^s}(x_k^s) - \nabla f(x_k^s)\|^2\right] \leq 4L[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)].$$

The proof of this lemma is included in APPENDIX A. Lemma 2 provides the upper bound on the expected variance of the variance reduced gradient estimator in (5a), i.e., the SVRG estimator independently introduced in [1], [35]. For Algorithm 3 with *Option I* and a fixed learning rate, we give the following key result for our analysis.

**Lemma 3** (Option I and smooth objectives). *Let $\beta = 1/(L\eta)$. If each $f_i(\cdot)$ is convex and $L$-smooth, then the following inequality holds for all $s = 1, 2, \ldots, S$,*

$$\begin{aligned}
&\left(1 - \frac{2}{\beta - 1}\right)\mathbb{E}[f(\widetilde{x}^s) - f(x^*)] + \frac{1}{m-1}\mathbb{E}[f(x_m^s) - f(x^*)] \\
\leq &\frac{2m}{(\beta-1)(m-1)}\mathbb{E}\left[f(\widetilde{x}^{s-1}) - f(x^*)\right] + \frac{2}{(\beta-1)(m-1)}\mathbb{E}[f(x_0^s) - f(x^*)] + \frac{L\beta}{2(m-1)}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

*Proof.* In order to simplify notation, the stochastic gradient estimator is defined as: $v_k^s := \nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1}) + \nabla f(\widetilde{x}^{s-1})$. Since each component function $f_i(x)$ is $L$-smooth, which implies that the gradient of the average function $f(x)$ is also $L$-smooth, i.e., for all $x, y \in \mathbb{R}^d$,

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|,$$

whose equivalent form is

$$f(y) \leq f(x) + \langle \nabla f(x),\ y - x\rangle + \frac{L}{2}\|y - x\|^2.$$

Applying the above smoothness inequality, we have

$$\begin{aligned}
f(x_{k+1}^s) &\leq f(x_k^s) + \left\langle \nabla f(x_k^s),\ x_{k+1}^s - x_k^s\right\rangle + \frac{L}{2}\left\|x_{k+1}^s - x_k^s\right\|^2 \\
&= f(x_k^s) + \left\langle \nabla f(x_k^s),\ x_{k+1}^s - x_k^s\right\rangle + \frac{L\beta}{2}\left\|x_{k+1}^s - x_k^s\right\|^2 - \frac{L(\beta-1)}{2}\left\|x_{k+1}^s - x_k^s\right\|^2 \\
&= f(x_k^s) + \left\langle v_k^s,\ x_{k+1}^s - x_k^s\right\rangle + \frac{L\beta}{2}\left\|x_{k+1}^s - x_k^s\right\|^2 \\
&\quad + \left\langle \nabla f(x_k^s) - v_k^s,\ x_{k+1}^s - x_k^s\right\rangle - \frac{L(\beta-1)}{2}\left\|x_{k+1}^s - x_k^s\right\|^2,
\end{aligned} \tag{15}$$

where $\beta = 1/(L\eta) > 3$ is a constant. Using Lemma 2, then we get

$$\begin{aligned}
&\mathbb{E}\left[\left\langle \nabla f(x_k^s) - v_k^s,\ x_{k+1}^s - x_k^s\right\rangle - \frac{L(\beta-1)}{2}\left\|x_{k+1}^s - x_k^s\right\|^2\right] \\
\leq &\mathbb{E}\left[\frac{1}{2L(\beta-1)}\|\nabla f(x_k^s) - v_k^s\|^2 + \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2 - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2\right] \\
\leq &\frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right],
\end{aligned} \tag{16}$$

where the first inequality holds due to the Young's inequality (i.e., $y^T z \leq \|y\|^2/(2\gamma) + \gamma\|z\|^2/2$ for all $\gamma > 0$ and $y, z \in \mathbb{R}^d$), and the second inequality follows from Lemma 2.

Substituting the inequality in (16) into the inequality in (15), and taking the expectation with respect to the random choice $i_k^s$, we have

$$\begin{aligned}
&\mathbb{E}[f(x_{k+1}^s)] \\
\leq &\mathbb{E}[f(x_k^s)] + \mathbb{E}\left[\left\langle v_k^s,\ x_{k+1}^s - x_k^s\right\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2\right] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right] \\
\leq &\mathbb{E}[f(x_k^s)] + \mathbb{E}\left[\left\langle v_k^s,\ x^* - x_k^s\right\rangle + \frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right] \\
\leq &\mathbb{E}[f(x_k^s)] + \mathbb{E}[\langle \nabla f(x_k^s),\ x^* - x_k^s\rangle] + \mathbb{E}\left[\frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right] \\
\leq &f(x^*) + \frac{L\beta}{2}\mathbb{E}[(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right].
\end{aligned}$$

Here, the first inequality holds due to the inequality in (15) and the inequality in (16); the second inequality follows from Lemma 1 with $\hat{z} = x_{k+1}^s$, $z = x^*$, $z_0 = x_k^s$, $\tau = L\beta = 1/\eta$, and $r(z) := \langle v_k^s, z - x_k^s\rangle$; the third inequality holds due to the fact that $\mathbb{E}[v_k^s] = \nabla f(x_k^s)$; and the last inequality follows from the convexity of the smooth function $f(\cdot)$, i.e., $f(x_k^s) + \langle \nabla f(x_k^s),\, x^* - x_k^s\rangle \le f(x^*)$. The above inequality can be rewritten as follows:

$$\mathbb{E}[f(x_{k+1}^s)] - f(x^*) \le \frac{2}{\beta - 1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right] + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2\right].$$

Summing the above inequality over $k = 0, 1, \ldots, m-1$, then

$$\sum_{k=1}^{m}\mathbb{E}[f(x_k^s) - f(x^*)] \le \sum_{k=1}^{m}\left\{\frac{2}{\beta - 1}\left[f(x_{k-1}^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right] + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_{k-1}^s\|^2 - \|x^* - x_k^s\|^2\right]\right\}. \tag{17}$$

Due to the setting of $\widetilde{x}^s = \frac{1}{m-1}\sum_{k=1}^{m-1} x_k^s$ in Option I, and the convexity of $f(\cdot)$, then we have

$$f(\widetilde{x}^s) \le \frac{1}{m-1}\sum_{k=1}^{m-1} f(x_k^s). \tag{18}$$

The left and right hand sides of the inequality in (17) can be rewritten as follows:

$$\sum_{k=1}^{m}\mathbb{E}[f(x_k^s) - f(x^*)] = \sum_{k=1}^{m-1}\mathbb{E}[f(x_k^s) - f(x^*)] + \mathbb{E}[f(x_m^s) - f(x^*)],$$

$$\begin{aligned}
&\sum_{k=1}^{m}\left\{\frac{2}{\beta - 1}\left[f(x_{k-1}^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right] + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_{k-1}^s\|^2 - \|x^* - x_k^s\|^2\right]\right\} \\
=&\frac{2}{\beta - 1}\sum_{k=1}^{m-1}[f(x_k^s) - f(x^*)] + \frac{2}{\beta - 1}\left\{f(x_0^s) - f(x^*) + m[f(\widetilde{x}^{s-1}) - f(x^*)]\right\} + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

Subtracting $\frac{2}{\beta-1}\sum_{k=1}^{m-1}[f(x_k^s) - f(x^*)]$ from both sides of the inequality in (17), then we obtain

$$\begin{aligned}
&\left(1 - \frac{2}{\beta - 1}\right)\sum_{k=1}^{m-1}\mathbb{E}[f(x_k^s) - f(x^*)] + \mathbb{E}[f(x_m^s) - f(x^*)] \\
\le&\frac{2}{\beta - 1}\mathbb{E}\left\{f(x_0^s) - f(x^*) + m[f(\widetilde{x}^{s-1}) - f(x^*)]\right\} + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

Applying the inequality in (18), we have

$$\begin{aligned}
&\left(1 - \frac{2}{\beta - 1}\right)(m-1)\mathbb{E}[f(\widetilde{x}^s) - f(x^*)] + \mathbb{E}[f(x_m^s) - f(x^*)] \\
\le&\left(1 - \frac{2}{\beta - 1}\right)\sum_{k=1}^{m-1}\mathbb{E}[f(x_k^s) - f(x^*)] + \mathbb{E}[f(x_m^s) - f(x^*)] \\
\le&\frac{2}{\beta - 1}\mathbb{E}\left\{f(x_0^s) - f(x^*) + m[f(\widetilde{x}^{s-1}) - f(x^*)]\right\} + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

Dividing both sides of the above inequality by $(m-1)$, we arrive at

$$\begin{aligned}
&\left(1 - \frac{2}{\beta - 1}\right)\mathbb{E}[f(\widetilde{x}^s) - f(x^*)] + \frac{1}{m-1}\mathbb{E}[f(x_m^s) - f(x^*)] \\
\le&\frac{2}{(\beta-1)(m-1)}\mathbb{E}[f(x_0^s) - f(x^*)] + \frac{2m}{(\beta-1)(m-1)}\mathbb{E}\left[f(\widetilde{x}^{s-1}) - f(x^*)\right] + \frac{L\beta}{2(m-1)}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

This completes the proof. □

The first main result is the following theorem, which provides the convergence guarantee of VR-SGD with Option I for solving smooth and general convex minimization problems.

**Theorem 1** (Option I and smooth objectives). *Suppose Assumption 1 holds. Then the following inequality holds*

$$\mathbb{E}\left[f(\widehat{x}^S)\right] - f(x^*) \le \frac{2(m+1)}{[(\beta-1)(m-1) - 4m + 2]S}[f(\widetilde{x}^0) - f(x^*)] + \frac{\beta(\beta-1)L}{2[(\beta-1)(m-1) - 4m + 2]S}\|\widetilde{x}^0 - x^*\|^2,$$

*where $\widehat{x}^S = \widetilde{x}^S$ if $f(\widetilde{x}^S) \le f(\overline{x}^S)$, and $\overline{x}^S = \frac{1}{S}\sum_{s=1}^{S}\widetilde{x}^s$. Otherwise, $\widehat{x}^S = \overline{x}^S$.*

*Proof.* Since $2/(\beta-1)<1$, it is easy to verify that

$$\frac{2}{(\beta-1)(m-1)}\left\{\mathbb{E}[f(x_m^s)]-f(x^*)\right\}\leq\frac{1}{m-1}\left\{\mathbb{E}[f(x_m^s)]-f(x^*)\right\}. \tag{19}$$

Applying the above inequality and Lemma 3, we have

$$\begin{aligned}
&\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]+\frac{2}{(\beta-1)(m-1)}\mathbb{E}[f(x_m^s)-f(x^*)]\\
\leq&\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]+\frac{1}{m-1}\mathbb{E}[f(x_m^s)-f(x^*)]\\
\leq&\frac{2}{(\beta-1)(m-1)}\mathbb{E}[f(x_0^s)-f(x^*)]+\frac{2m}{(\beta-1)(m-1)}\mathbb{E}\big[f(\widetilde{x}^{s-1})-f(x^*)\big]+\frac{L\beta}{2(m-1)}\mathbb{E}\big[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2\big].
\end{aligned}$$

Summing the above inequality over $s=1,2,\ldots,S$, taking expectation with respect to the history of random variables $i_k^s$, and using the setting of $x_0^{s+1}=x_m^s$, we obtain

$$\begin{aligned}
&\sum_{s=1}^{S}\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]\\
\leq&\sum_{s=1}^{S}\left\{\frac{2}{(\beta-1)(m-1)}\mathbb{E}[f(x_0^s)-f(x^*)-(f(x_m^s)-f(x^*))]+\frac{2m}{(\beta-1)(m-1)}\mathbb{E}[f(\widetilde{x}^{s-1})-f(x^*)]\right\}\\
&+\frac{L\beta}{2(m-1)}\sum_{s=1}^{S}\mathbb{E}\big[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2\big].
\end{aligned}$$

Subtracting $\frac{2m}{(\beta-1)(m-1)}\sum_{s=1}^{S-1}[f(\widetilde{x}^s)-f(x^*)]$ from both sides of the above inequality, we have

$$\begin{aligned}
&\frac{2m}{(\beta-1)(m-1)}\mathbb{E}[f(\widetilde{x}^S)-f(x^*)]+\sum_{s=1}^{S}\left(1-\frac{4}{\beta-1}-\frac{2}{(\beta-1)(m-1)}\right)\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]\\
\leq&\frac{2}{(\beta-1)(m-1)}\mathbb{E}\Big[f(x_0^1)-f(x^*)-(f(x_m^S)-f(x^*))\Big]+\frac{2m}{(\beta-1)(m-1)}\mathbb{E}[f(\widetilde{x}^0)-f(x^*)]\\
&+\frac{L\beta}{2(m-1)}\mathbb{E}\Big[\|x^*-x_0^1\|^2-\|x^*-x_m^S\|^2\Big].
\end{aligned}$$

Dividing both sides of the above inequality by $S$, and using the setting of $\widetilde{x}^0=x_0^1$, we arrive at

$$\begin{aligned}
&\frac{1}{S}\left(1-\frac{4}{\beta-1}-\frac{2}{(\beta-1)(m-1)}\right)\sum_{s=1}^{S}\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]\\
\leq&\frac{2m}{(\beta-1)(m-1)S}\mathbb{E}[f(\widetilde{x}^S)-f(x^*)]+\frac{1}{S}\left(1-\frac{4}{\beta-1}-\frac{2}{(\beta-1)(m-1)}\right)\sum_{s=1}^{S}\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]\\
\leq&\frac{2}{(\beta-1)(m-1)S}[f(x_0^1)-f(x^*)]+\frac{2m}{(\beta-1)(m-1)S}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta}{2(m-1)S}\|x^*-x_0^1\|^2\\
=&\frac{2(m+1)}{(\beta-1)(m-1)S}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta}{2(m-1)S}\|\widetilde{x}^0-x^*\|^2,
\end{aligned}$$

where the first inequality holds due to the fact that $f(\widetilde{x}^S)-f(x^*)\geq 0$; the second inequality holds due to the facts that $f(x_m^S)-f(x^*)\geq 0$ and $\|x^*-x_m^S\|^2\geq 0$; and the last equality follows from the setting of $\widetilde{x}^0=x_0^1$.

Due to the definition of $\overline{x}^S$ and the convexity of $f(\cdot)$, we have $f(\overline{x}^S)\leq\frac{1}{S}\sum_{s=1}^{S}f(\widetilde{x}^s)$, and therefore the above inequality becomes:

$$\begin{aligned}
&\left(1-\frac{4}{\beta-1}-\frac{2}{(\beta-1)(m-1)}\right)\mathbb{E}\Big[f(\overline{x}^S)-f(x^*)\Big]\\
\leq&\frac{2(m+1)}{(\beta-1)(m-1)S}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta}{2(m-1)S}\|\widetilde{x}^0-x^*\|^2.
\end{aligned} \tag{20}$$

Dividing both sides of the above inequality by $c_1=1-\frac{4}{\beta-1}-\frac{2}{(\beta-1)(m-1)}>0$, we have

$$\mathbb{E}\Big[f(\overline{x}^S)\Big]-f(x^*)\leq\frac{2(m+1)}{c_1(\beta-1)(m-1)S}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta}{2c_1(m-1)S}\|\widetilde{x}^0-x^*\|^2.$$

Due to the setting for the output of Algorithm 3, $\widehat{x}^S = \widetilde{x}^S$ if $f(\widetilde{x}^S) \leq f(\overline{x}^S)$. Then

$$\mathbb{E}\Big[f(\widehat{x}^S)\Big] - f(x^*) \leq \mathbb{E}\Big[f(\overline{x}^S)\Big] - f(x^*) \leq \frac{2(m+1)}{c_1(\beta-1)(m-1)S}[f(\widetilde{x}^0) - f(x^*)] + \frac{L\beta}{2c_1(m-1)S}\|\widetilde{x}^0 - x^*\|^2.$$

Alternatively, when $f(\widetilde{x}^S) \geq f(\overline{x}^S)$, let $\widehat{x}^S = \overline{x}^S$, and the above inequality still holds. This completes the proof. □

From Lemma 3, Theorem 1, and their proofs, one can see that our convergence analysis is very different from those of existing stochastic methods, such as SVRG [1] and Prox-SVRG [2]. For Algorithm 3 with *Option II* and a fixed learning rate, we give the following key lemma for our convergence analysis.

**Lemma 4** (Option II and smooth objectives). *If each $f_i(\cdot)$ is convex and $L$-smooth, then the following inequality holds for all $s = 1, 2, \ldots, S$,*

$$\begin{aligned}&\left(1 - \frac{2}{\beta-1}\right)\mathbb{E}[f(\widetilde{x}^s) - f(x^*)] + \frac{2}{(\beta-1)m}\mathbb{E}[f(x_m^s) - f(x^*)] \\ \leq\ & \frac{2}{(\beta-1)}\mathbb{E}\big[f(\widetilde{x}^{s-1}) - f(x^*)\big] + \frac{2}{(\beta-1)m}\mathbb{E}[f(x_0^s) - f(x^*)] + \frac{L\beta}{2m}\mathbb{E}\big[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\big].\end{aligned}$$

This lemma is a slight generalization of Lemma 3, and we give the proof in APPENDIX B for completeness.

**Theorem 2** (Option II and smooth objectives). *If each $f_i(\cdot)$ is convex and $L$-smooth, then the following inequality holds*

$$\mathbb{E}\Big[f(\widehat{x}^S)\Big] - f(x^*) \leq \frac{2(m+1)}{mS(\beta-5)}[f(\widetilde{x}^0) - f(x^*)] + \frac{L\beta(\beta-1)}{2mS(\beta-5)}\|\widetilde{x}^0 - x^*\|^2,$$

*where $\widehat{x}^S = \widetilde{x}^S$ if $f(\widetilde{x}^S) \leq f(\overline{x}^S)$, and $\overline{x}^S = \frac{1}{S}\sum_{s=1}^S \widetilde{x}^s$. Otherwise, $\widehat{x}^S = \overline{x}^S$.*

The proof of this theorem can be found in APPENDIX C. Clearly, Theorems 1 and 2 show that VR-SGD with Option I or Option II attains a sub-linear convergence rate for smooth general convex objective functions. This means that VR-SGD is guaranteed to have a similar convergence rate with other variance reduced stochastic methods such as SAGA [30], and a slower theoretical rate than accelerated methods such as Katyusha [3]. Nevertheless, VR-SGD usually converges much faster than the best known stochastic method, Katyusha [3] in practice (see Section 5.3 for details).

#### 4.1.2 Non-Smooth Objectives

Next, we provide the convergence guarantee of Problem (1) when the objective function $F(x)$ is *non-smooth* (i.e., the regularizer $g(x)$ is non-smooth) and non-strongly convex. We first give the following lemma.

**Lemma 5** (Variance bound of non-smooth objectives). *Let $x^*$ be the optimal solution of Problem (1). If each $f_i(\cdot)$ is convex and $L$-smooth, and $\widetilde{\nabla} f_{i_k^s}(x_k^s) := \nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1}) + \nabla f(\widetilde{x}^{s-1})$, then the following inequality holds*

$$\mathbb{E}\Big[\|\widetilde{\nabla} f_{i_k^s}(x_k^s) - \nabla f(x_k^s)\|^2\Big] \leq 4L[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)].$$

This lemma can be viewed as a generalization of Lemma 2, and is essentially identical to Corollary 3.5 in [2] and Lemma 1 in [45], and hence its proof is omitted. For Algorithm 3 with *Option II* and a fixed learning rate, we give the following results.

**Lemma 6** (Option II and non-smooth objectives). *If each $f_i(\cdot)$ is convex and $L$-smooth, then the following inequality holds for all $s = 1, 2, \ldots, S$,*

$$\begin{aligned}&\left(1 - \frac{2}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s) - F(x^*)] + \frac{2}{(\beta-1)m}\mathbb{E}[F(x_m^s) - F(x^*)] \\ \leq\ & \frac{2}{(\beta-1)m}\mathbb{E}[F(x_0^s) - F(x^*)] + \frac{2}{\beta-1}\mathbb{E}\big[F(\widetilde{x}^{s-1}) - F(x^*)\big] + \frac{L\beta}{2m}\mathbb{E}\big[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\big].\end{aligned} \tag{21}$$

The proof of this lemma can be found in APPENDIX D. Using the above lemma, we give the following convergence result for VR-SGD.

**Theorem 3** (Option II and non-smooth objectives). *Suppose Assumption 1 holds. Then the following inequality holds*

$$\mathbb{E}\Big[F(\widehat{x}^S)\Big] - F(x^*) \leq \frac{2(m+1)}{(\beta-5)mS}[F(\widetilde{x}^0) - F(x^*)] + \frac{\beta(\beta-1)L}{2(\beta-5)mS}\|\widetilde{x}^0 - x^*\|^2.$$

The proof of this theorem is included in APPENDIX E. Similarly, for Algorithm 3 with *Option I* and a fixed learning rate, we give the following results.

**Corollary 1** (Option I and non-smooth objectives). *If each $f_i(\cdot)$ is convex and $L$-smooth, then the following inequality holds for all $s = 1, 2, \ldots, S$,*

$$\begin{aligned}&\left(1 - \frac{2}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s) - F(x^*)] + \frac{1}{m-1}\mathbb{E}[F(x_m^s) - F(x^*)]\\ \leq &\frac{2m}{(\beta-1)(m-1)}\mathbb{E}\big[F(\widetilde{x}^{s-1}) - F(x^*)\big] + \frac{2}{(\beta-1)(m-1)}\mathbb{E}[F(x_0^s) - F(x^*)] + \frac{L\beta}{2(m-1)}\mathbb{E}\big[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\big].\end{aligned}$$

**Corollary 2** (Option I and non-smooth objectives). *Suppose Assumption 1 holds. Then the following inequality holds*

$$\begin{aligned}&\mathbb{E}\Big[F(\widehat{x}^S)\Big] - F(x^*)\\ \leq &\frac{2(m+1)}{[(\beta-1)(m-1)-4m+2]S}[F(\widetilde{x}^0) - F(x^*)] + \frac{\beta(\beta-1)L}{2[(\beta-1)(m-1)-4m+2]S}\|\widetilde{x}^0 - x^*\|^2,\end{aligned}$$

*where $\widehat{x}^S = \widetilde{x}^S$ if $F(\widetilde{x}^S) \leq F(\overline{x}^S)$, and $\overline{x}^S = \frac{1}{S}\sum_{s=1}^S \widetilde{x}^s$. Otherwise, $\widehat{x}^S = \overline{x}^S$.*

Corollaries 1 and 2 can be viewed as the generalizations of Lemma 3 and Theorem 1, respectively, and hence their proofs are omitted. Obviously, Theorems 3 and Corollary 2 show that VR-SGD with Option I or Option II attains a sub-linear convergence rate for non-smooth general convex objective functions.

#### 4.1.3 Mini-Batch Settings

The upper bound on the variance of the stochastic gradient estimator $\widetilde{\nabla} f_{i_k^s}(x_k^s)$ in Lemma 5 is extended to the *mini-batch setting* as follows.

**Lemma 7** (Variance bound of mini-batch). *If each $f_i(\cdot)$ is convex and $L$-smooth, and $\widetilde{\nabla} f_{I_k^s}(x_k^s) := \frac{1}{b}\sum_{i\in I_k^s}\big[\nabla f_i(x_k^s) - \nabla f_i(\widetilde{x}^{s-1})\big] + \nabla f(\widetilde{x}^{s-1})$, then the following inequality holds*

$$\mathbb{E}\Big[\|\widetilde{\nabla} f_{I_k^s}(x_k^s) - \nabla f(x_k^s)\|^2\Big] \leq 4L\delta(b)[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)],$$

*where $\delta(b) = (n-b)/[(n-1)b]$.*

It is not hard to verify that $0 \leq \delta(b) \leq 1$. This lemma is essentially identical to Theorem 4 in [36], and hence its proof is omitted. Based on the variance upper bound in Lemma 7, we further analyze the convergence properties of VR-SGD for the mini-batch setting. Lemma 6 is first extended to the mini-batch setting as follows.

**Lemma 8** (Mini-batch). *Using the same notation as in Lemma 7, we have*

$$\begin{aligned}&\frac{2\delta(b)}{(\beta-1)m}\{\mathbb{E}[F(x_m^s)] - F(x^*)\} + \left(1 - \frac{2\delta(b)}{\beta-1}\right)\{\mathbb{E}[F(\widetilde{x}^s)] - F(x^*)\}\\ \leq &\frac{2\delta(b)}{(\beta-1)}\left[F(\widetilde{x}^{s-1}) - F(x^*)\right] + \frac{2\delta(b)}{(\beta-1)m}\left[F(x_0^s) - F(x^*)\right] + \frac{L\beta}{2m}\mathbb{E}\big[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\big].\end{aligned}$$

*Proof.* In order to simplify notation, the stochastic gradient estimator of mini-batch is defined as:

$$v_k^s := \frac{1}{b}\sum_{i\in I_k^s}\big[\nabla f_i(x_k^s) - \nabla f_i(\widetilde{x}^{s-1})\big] + \nabla f(\widetilde{x}^{s-1}).$$

$$\begin{aligned}F(x_{k+1}^s) &\leq g(x_{k+1}^s) + f(x_k^s) + \langle \nabla f(x_k^s),\, x_{k+1}^s - x_k^s\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2 - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2\\ &= g(x_{k+1}^s) + f(x_k^s) + \langle v_k^s,\, x_{k+1}^s - x_k^s\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2\\ &\quad + \langle \nabla f(x_k^s) - v_k^s,\, x_{k+1}^s - x_k^s\rangle - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2.\end{aligned} \tag{22}$$

Using Lemma 7, then we obtain

$$\begin{aligned}
&\mathbb{E}\left[\langle \nabla f(x_k^s) - v_k^s,\, x_{k+1}^s - x_k^s\rangle - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2\right] \\
\leq &\mathbb{E}\left[\frac{1}{2L(\beta-1)}\|\nabla f(x_k^s) - v_k^s\|^2 + \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2 - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2\right] \\
\leq &\frac{2\delta(b)}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right],
\end{aligned} \tag{23}$$

where the first inequality holds due to the Young's inequality, and the second inequality follows from Lemma 7. Substituting the inequality (23) into the inequality (22), and taking the expectation over the random mini-batch set $I_k^s$, we have

$$\begin{aligned}
&\mathbb{E}[F(x_{k+1}^s)] \\
\leq &\mathbb{E}[g(x_{k+1}^s)] + \mathbb{E}[f(x_k^s)] + \mathbb{E}\left[\langle v_k^s,\, x_{k+1}^s - x_k^s\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2\right] + \frac{2\delta(b)}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] \\
\leq &g(x^*) + \mathbb{E}[f(x_k^s)] + \mathbb{E}\left[\langle v_k^s,\, x^* - x_k^s\rangle + \frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2\delta(b)}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] \\
\leq &g(x^*) + f(x^*) + \mathbb{E}\left[\frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2\delta(b)}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] \\
= &F(x^*) + \frac{L\beta}{2}\mathbb{E}\left[(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2\delta(b)}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right],
\end{aligned}$$

where the second inequality holds from Lemma 1. Then the above inequality is rewritten as follows:

$$\mathbb{E}\left[F(x_{k+1}^s)\right] - F(x^*) \leq \frac{2\delta(b)}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2\right]. \tag{24}$$

Summing the above inequality over $k = 0, 1, \cdots, (m-1)$, then

$$\begin{aligned}
&\sum_{k=0}^{m-1}\left\{\mathbb{E}[F(x_{k+1}^s)] - F(x^*)\right\} \\
\leq &\sum_{k=0}^{m-1}\left\{\frac{2\delta(b)}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2\right]\right\}.
\end{aligned}$$

Since $\widetilde{x}^s = \frac{1}{m}\sum_{k=1}^m x_k^s$, we have $F(\widetilde{x}^s) \leq \frac{1}{m}\sum_{k=1}^m F(x_k^s)$, and

$$\begin{aligned}
&\frac{2\delta(b)}{(\beta-1)m}\left\{\mathbb{E}[F(x_m^s)] - F(x^*)\right\} + \left(1 - \frac{2\delta(b)}{\beta-1}\right)\left\{\mathbb{E}[F(\widetilde{x}^s)] - F(x^*)\right\} \\
\leq &\frac{2\delta(b)}{(\beta-1)}\left[F(\widetilde{x}^{s-1}) - F(x^*)\right] + \frac{2\delta(b)}{(\beta-1)m}\left[F(x_0^s) - F(x^*)\right] + \frac{L\beta}{2m}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

This completes the proof. □

Similar to Lemma 8, we can also extend Lemma 3 and Corollary 1 to the mini-batch setting.

**Theorem 4** (Mini-batch). *If each $f_i(\cdot)$ is convex and $L$-smooth, then the following inequality holds*

$$\mathbb{E}\left[F(\widehat{x}^S)\right] - F(x^*) \leq \frac{2\delta(b)(m+1)}{(\beta-1-4\delta(b))mS}\mathbb{E}\left[F(\widetilde{x}^0) - F(x^*)\right] + \frac{L\beta(\beta-1)}{2(\beta-1-4\delta(b))mS}\mathbb{E}\left[\|x^* - \widetilde{x}^0\|^2\right]. \tag{25}$$

*Proof.* Using Lemma 8, we have

$$\begin{aligned}
&\left(1 - \frac{2\delta(b)}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s) - F(x^*)] \\
\leq &\frac{2\delta(b)}{(\beta-1)}\left[F(\widetilde{x}^{s-1}) - F(x^*)\right] + \frac{2\delta(b)}{(\beta-1)m}\left[F(x_0^s) - F(x^*)\right] + \frac{L\beta}{2m}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

Summing the above inequality over $s = 1, 2, \ldots, S$, taking expectation over whole history of $I_k^s$, and using $x_0^{s+1} = x_m^s$, we obtain

$$\begin{aligned}
&\sum_{s=1}^S\left(1 - \frac{2\delta(b)}{\beta-1}\right)\mathbb{E}\left[F(\widetilde{x}^s) - F(x^*)\right] \\
\leq &\sum_{s=1}^S\frac{2\delta(b)}{(\beta-1)}\left[F(\widetilde{x}^{s-1}) - F(x^*)\right] + \sum_{s=1}^S\frac{2\delta(b)}{(\beta-1)m}\mathbb{E}[F(x_0^s) - F(x_m^s)] + \frac{L\beta}{2m}\sum_{s=1}^S\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

Subtracting $\frac{2\delta(b)}{\beta-1}\sum_{s=1}^{S-1}\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]$ from both sides of the above inequality, we have

$$\begin{aligned}&\frac{2\delta(b)}{\beta-1}\mathbb{E}\Big[F(\widetilde{x}^S)-F(x^*)\Big]+\sum_{s=1}^{S}\left(1-\frac{4\delta(b)}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]\\ \leq&\frac{2\delta(b)}{(\beta-1)}\left[F(\widetilde{x}^0)-F(x^*)\right]+\frac{2\delta(b)}{(\beta-1)m}\mathbb{E}\Big[F(x_0^1)-F(x_m^S)\Big]+\frac{L\beta}{2m}\mathbb{E}\Big[\|x^*-x_0^1\|^2-\|x^*-x_m^S\|^2\Big]\end{aligned}$$

Dividing both sides of the above inequality by $S$ and using $\mathbb{E}[F(\overline{x})]\leq\frac{1}{S}\sum_{s=1}^{S}F(\widetilde{x}^s)$, we arrive at

$$\begin{aligned}&\frac{2\delta(b)}{(\beta-1)S}\mathbb{E}[F(\widetilde{x}^S)-F(x^*)]+\left(1-\frac{4\delta(b)}{\beta-1}\right)\mathbb{E}[F(\overline{x})-F(x^*)]\\ \leq&\frac{2\delta(b)}{(\beta-1)S}\left[F(\widetilde{x}^0)-F(x^*)\right]+\frac{2\delta(b)}{(\beta-1)mS}\mathbb{E}\Big[F(x_0^1)-F(x_m^S)\Big]+\frac{L\beta}{2mS}\mathbb{E}\Big[\|x^*-x_0^1\|^2-\|x^*-x_m^S\|^2\Big].\end{aligned}$$

Subtracting $\frac{2\delta(b)}{(\beta-1)S}\mathbb{E}[F(\widetilde{x}^S)-F(x^*)]$ from both sides of the above inequality, we have

$$\begin{aligned}&\left(1-\frac{4\delta(b)}{\beta-1}\right)\mathbb{E}[F(\overline{x})-F(x^*)]\\ \leq&\frac{2\delta(b)}{(\beta-1)S}\mathbb{E}\Big[F(\widetilde{x}^0)-F(\widetilde{x}^S)\Big]+\frac{2\delta(b)}{(\beta-1)mS}\mathbb{E}\Big[F(x_0^1)-F(x_m^S)\Big]+\frac{L\beta}{2mS}\mathbb{E}\Big[\|x^*-x_0^1\|^2-\|x^*-x_m^S\|^2\Big].\end{aligned}$$

Dividing both sides of the above inequality by $(1-\frac{4\delta(b)}{\beta-1})>0$, we arrive at

$$\begin{aligned}&\mathbb{E}[F(\overline{x})]-F(x^*)\\ \leq&\frac{2\delta(b)}{(\beta-1-4\delta(b))S}\mathbb{E}\Big[F(\widetilde{x}^0)-F(\widetilde{x}^S)\Big]+\frac{2\delta(b)}{(\beta-1-4\delta(b))mS}\mathbb{E}\Big[F(\widetilde{x}^0)-F(x_m^S)\Big]+\frac{L\beta(\beta-1)}{2(\beta-1-4\delta(b))mS}\mathbb{E}\big[\|x^*-\widetilde{x}^0\|^2\big]\\ \leq&\frac{2\delta(b)}{(\beta-1-4\delta(b))S}\mathbb{E}\big[F(\widetilde{x}^0)-F(x^*)\big]+\frac{2\delta(b)}{(\beta-1-4\delta(b))mS}\mathbb{E}\big[F(\widetilde{x}^0)-F(x^*)\big]+\frac{L\beta(\beta-1)}{2(\beta-1-4\delta(b))mS}\mathbb{E}\big[\|x^*-\widetilde{x}^0\|^2\big]\\ =&\frac{2\delta(b)(m+1)}{(\beta-1-4\delta(b))mS}\mathbb{E}\big[F(\widetilde{x}^0)-F(x^*)\big]+\frac{L\beta(\beta-1)}{2(\beta-1-4\delta(b))mS}\mathbb{E}\big[\|x^*-\widetilde{x}^0\|^2\big].\end{aligned}$$

When $F(\widetilde{x}^S)\leq F(\overline{x}^S)$, then $\widehat{x}^S=\widetilde{x}^S$, and

$$\mathbb{E}\Big[F(\widehat{x}^S)\Big]-F(x^*)\leq\frac{2\delta(b)(m+1)}{(\beta-1-4\delta(b))mS}\mathbb{E}\big[F(\widetilde{x}^0)-F(x^*)\big]+\frac{L\beta(\beta-1)}{2(\beta-1-4\delta(b))mS}\mathbb{E}\big[\|x^*-\widetilde{x}^0\|^2\big].$$

Alternatively, if $F(\widetilde{x}^S)\geq F(\overline{x}^S)$, then $\widehat{x}^S=\overline{x}^S$, and the above inequality still holds. This completes the proof. □

From Theorem 4, one can see that when $b=n$ (i.e., the batch setting), we have $\delta(n)=0$, and the first term on the right-hand side of (25) diminishes. In other words, our VR-SGD method degenerates to the deterministic method with the convergence rate of $\mathcal{O}(1/T)$. Furthermore, when $b=1$, we have $\delta(1)=1$, and then Theorem 4 degenerates to Theorem 3.

### 4.2 Convergence Properties for Strongly Convex Problems

In this part, we analyze the convergence properties of VR-SGD (i.e., Algorithm 2) for solving the *strongly convex* objective function (1). According to the above analysis, one can see that if $F(\cdot)$ is convex, and each convex component function $f_i(\cdot)$ is $L$-smooth, both Algorithms 2 and 3 converge to the optimal solution. In the following, we provide stronger convergence rate guarantees for VR-SGD under the strongly convex condition. We first give the following assumption.

**Assumption 3.** *For all $s=1,2,\ldots,S$, the following inequality holds*

$$\mathbb{E}[F(x_0^s)-F(x^*)]\leq C\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]$$

*where $C>0$ is a constant*[4].

Similar to Algorithm 3, Algorithm 2 also has two different update rules for the two cases of smooth and non-smooth objective functions. We first give the following convergence result for Algorithm 2 with *Option I*.

4. This assumption shows the relationship of the gaps between the function values at the starting and snapshot points of each epoch and the optimal value of the objective function. In fact, both $\widetilde{x}^s$ and $x_m^s$ (i.e., $x_0^{s+1}$) converge to $x^*$, and thus $C$ is far less than $m$, i.e., $C\ll m$.

**Theorem 5** (Option I). *Suppose Assumptions 1, 2, and 3 hold, and $m$ is sufficiently large so that*

$$\rho_I := \frac{2L\eta(m+C)}{(m-1)(1-3L\eta)} + \frac{C(1-L\eta)}{\mu\eta(m-1)(1-3L\eta)} < 1.$$

*Then Algorithm 2 with Option I has the following geometric convergence in expectation:*

$$\mathbb{E}\left[F(\widetilde{x}^s) - F(x^*)\right] \leq \rho_I^s \left[F(\widetilde{x}^0) - F(x^*)\right].$$

*Proof.* Since each $f_i(\cdot)$ is convex and $L$-smooth, then Corollary 1 holds, which then implies

$$\begin{aligned}&\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]+\frac{1}{m-1}\mathbb{E}\big[F(x_0^{s+1})-F(x^*)\big]+\frac{L\beta}{2(m-1)}\mathbb{E}\big[\|x_0^{s+1}-x^*\|^2\big]\\ \leq&\frac{2m}{(m-1)(\beta-1)}\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]+\frac{2}{(m-1)(\beta-1)}\mathbb{E}[F(x_0^s)-F(x^*)]+\frac{L\beta}{2(m-1)}\mathbb{E}\big[\|x_0^s-x^*\|^2\big].\end{aligned} \tag{26}$$

Due to the strong convexity of $F(\cdot)$, we have $\|x_0^s - x^*\|^2 \leq (2/\mu)[F(x_0^s) - F(x^*)]$. Then the inequality in (26) can be rewritten as follows:

$$\begin{aligned}&\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]\\ \leq&\frac{2m}{(m-1)(\beta-1)}\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]+\left(\frac{2}{(m-1)(\beta-1)}+\frac{L\beta}{\mu(m-1)}\right)\mathbb{E}[F(x_0^s)-F(x^*)]\\ \leq&\frac{2m}{(m-1)(\beta-1)}\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]+\left(\frac{2C}{(m-1)(\beta-1)}+\frac{CL\beta}{\mu(m-1)}\right)\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]\\ =&\left(\frac{2(m+C)}{(m-1)(\beta-1)}+\frac{CL\beta}{\mu(m-1)}\right)\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]\end{aligned}$$

where the first inequality holds due to the fact that $\|x_0^s - x^*\|^2 \leq (2/\mu)[F(x_0^s) - F(x^*)]$, and the second inequality follows from Assumption 3. Dividing both sides of the above inequality by $[1-2/(\beta-1)]>0$ (that is, $\beta$ is required to be larger than 3) and using the definition of $\beta=1/(L\eta)$, we arrive at

$$\begin{aligned}\mathbb{E}[F(\widetilde{x}^s)-F(x^*)] \leq& \left(\frac{2(m+C)}{(m-1)(\beta-3)}+\frac{CL\beta(\beta-1)}{\mu(m-1)(\beta-3)}\right)\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]\\ =&\left(\frac{2(m+C)L\eta}{(m-1)(1-3L\eta)}+\frac{C(1-L\eta)}{\mu\eta(m-1)(1-3L\eta)}\right)\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big].\end{aligned}$$

This completes the proof. □

Although the learning rate $\eta$ in Theorem 5 needs to be less than $1/(3L)$, we can use much larger learning rates in practice, e.g., $\eta = 3/(7L)$. Similar to Theorem 5, we give the following convergence result for Algorithm 2 with *Option II*.

**Theorem 6** (Option II). *Suppose Assumptions 1, 2, and 3 hold, and $m$ is sufficiently large so that*

$$\rho_{II} := \frac{2L\eta(m+C)}{m(1-3L\eta)} + \frac{C(1-L\eta)}{m\mu\eta(1-3L\eta)} < 1.$$

*Then Algorithm 2 with Option II has the following geometric convergence in expectation:*

$$\mathbb{E}\left[F(\widetilde{x}^s) - F(x^*)\right] \leq \rho_{II}^s \left[F(\widetilde{x}^0) - F(x^*)\right].$$

The proof of Theorem 6 can be found in APPENDIX F. In addition, we can also provide the linear convergence guarantees of Algorithm 2 with Option I or Option II for solving smooth strongly convex objective functions. From all the results, one can see that VR-SGD attains a *linear* convergence rate for both *smooth* and *non-smooth* strongly convex minimization problems.

## 5 Experimental Results

In this section, we evaluate the performance of our VR-SGD method for solving various ERM problems, such as logistic regression, Lasso, and ridge regression, and compare its performance with several related stochastic variance reduced methods, including SVRG [1], Prox-SVRG [2], and Katyusha [3]. All the codes of VR-SGD and related methods can be downloaded from the author's website[5].

5. https://sites.google.com/site/fanhua217/publications

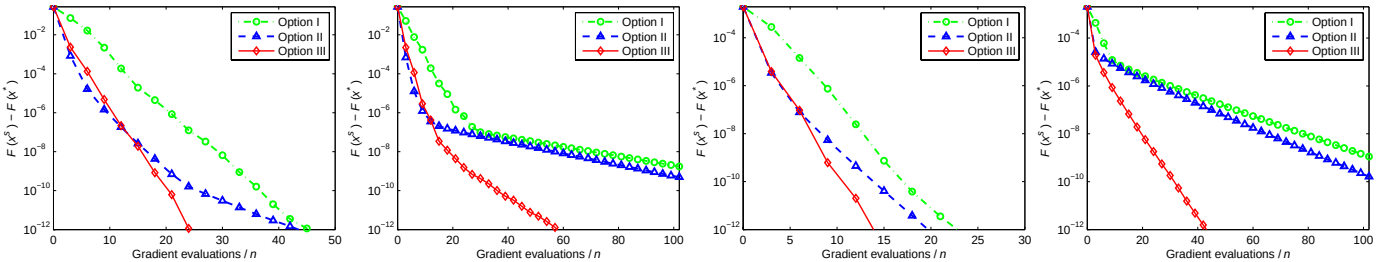


(a) Adult: $\lambda = 10^{-5}$ (left) and $\lambda = 10^{-6}$ (right)

(b) Protein: $\lambda = 10^{-5}$ (left) and $\lambda = 10^{-6}$ (right)

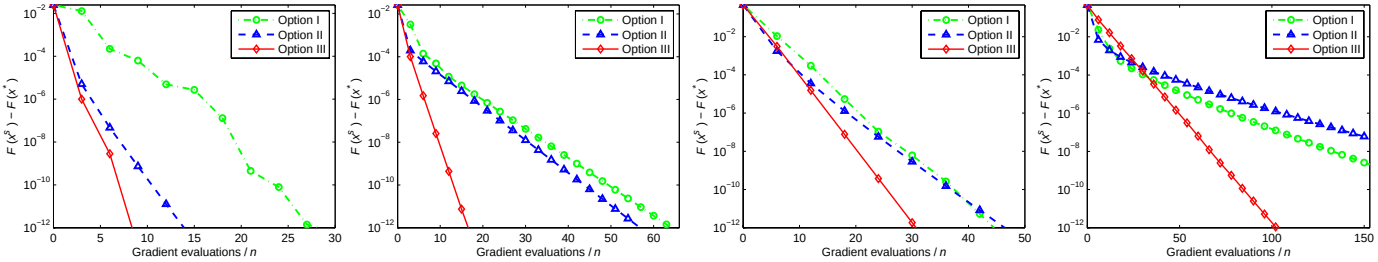


(c) Covtype: $\lambda = 10^{-5}$ (left) and $\lambda = 10^{-6}$ (right)

(d) Sido0: $\lambda = 10^{-4}$ (left) and $\lambda = 10^{-5}$ (right)

Fig. 3. Comparison of Options I, II, and III for solving ridge regression problems with the regularizer $(\lambda/2)\|\cdot\|^2$. In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes.

### 5.1 Experimental Setup

We used four publicly available data sets in the experiments: Adult (also called a9a), Covtype, Protein, and Sido0, as listed in Table 1. *It should be noted that each example of these date sets was normalized so that they have unit length as in [2], [34], which leads to the same upper bound on the Lipschitz constants $L_i$, i.e., $L = L_i$ for all $i = 1, \ldots, n$.* As suggested in [1], [2], [3], the epoch length is set to $m = 2n$ for the stochastic variance reduced methods, SVRG [1], Prox-SVRG [2], and Katyusha [3], as well as our VR-SGD method. Then the only parameter we have to tune by hand is the step size (or learning rate), $\eta$. Since Katyusha has a much higher per-iteration complexity than SVRG, Prox-SVRG, and VR-SGD, we compare their performance in terms of both the number of effective passes and running time (seconds), where computing a single full gradient or evaluating $n$ component gradients is considered as one effective pass over the data. For fair comparison, we implemented SVRG, Prox-SVRG, Katyusha, and VR-SGD in C++ with a Matlab interface, and performed all the experiments on a PC with an Intel i5-2400 CPU and 16GB RAM. In addition, we do not compare with other stochastic algorithms such as SAGA [30] and Catalyst [40], as they have been shown to be comparable or inferior to Katyusha [3].

TABLE 1
Summary of data sets used for our experiments.

| Data sets | Sizes $n$ | Dimensions $d$ | Sparsity |
|---|---|---|---|
| Adult | 32,562 | 123 | 11.28% |
| Protein | 145,751 | 74 | 99.21% |
| Covtype | 581,012 | 54 | 22.12% |
| Sido0 | 12,678 | 4,932 | 9.84% |

TABLE 2
The three choices of snapshot and starting points for stochastic optimization.

| Option I | Option II | Option III |
|---|---|---|
| $\widetilde{x}^{s+1} = x_m^s$ and $x_0^{s+1} = x_m^s$ | $\widetilde{x}^{s+1} = \frac{1}{m}\sum_{k=1}^m x_k^s$ and $x_0^{s+1} = \frac{1}{m}\sum_{k=1}^m x_k^s$ | $\widetilde{x}^{s+1} = \frac{1}{m-1}\sum_{k=1}^{m-1} x_k^s$ and $x_0^{s+1} = x_m^s$ |

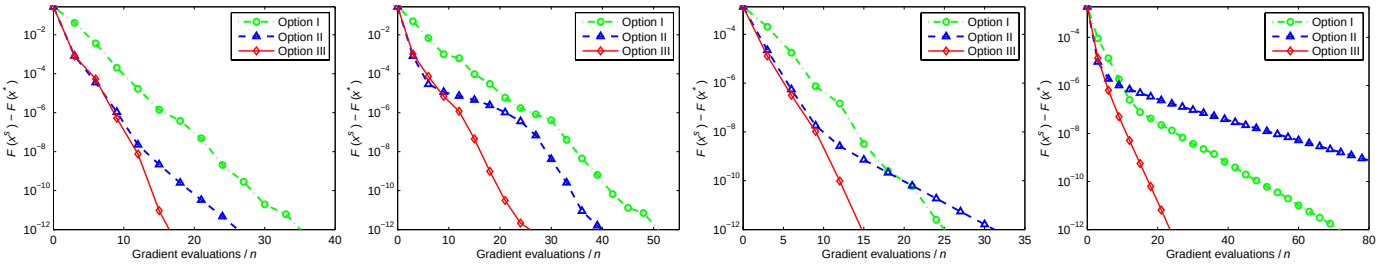


(a) Adult: $\lambda = 10^{-4}$ (left) and $\lambda = 10^{-5}$ (right)

(b) Protein: $\lambda = 10^{-4}$ (left) and $\lambda = 10^{-5}$ (right)

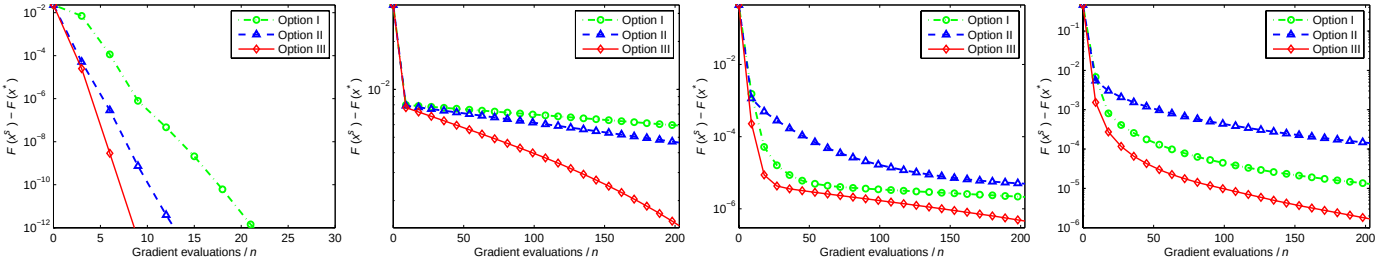


(c) Covtype: $\lambda = 10^{-4}$ (left) and $\lambda = 10^{-5}$ (right)

(d) Sido0: $\lambda = 10^{-4}$ (left) and $\lambda = 10^{-5}$ (right)

Fig. 4. Comparison of Options I, II, and III for solving Lasso problems with the regularizer $\lambda\|x\|_1$. In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes.

### 5.2 Different Choices for Snapshot and Starting Points

In the practical implementation of SVRG [1], both the snapshot point $\widetilde{x}^s$ and the starting point $x_0^{s+1}$ of each epoch are set to the last iterate $x_m^s$ of the previous epoch (i.e., Option I in Algorithm 1), while the two vectors in [2] (also suggested in [1]) are set to the average point of the previous epoch, $\frac{1}{m}\sum_{k=1}^{m} x_k^s$ (i.e., Option II in Algorithm 1). In contrast, $\widetilde{x}^s$ and $x_0^{s+1}$ in this paper are set to $\frac{1}{m-1}\sum_{k=1}^{m-1} x_k^s$ and $x_m^s$ (denoted by Option III, i.e., Option I[6] in Algorithms 2 and 3), respectively. Note that Johnson and Zhang [1] first presented the choices of Options I and II, while the setting of Option III is suggested in this paper. In the following, we compare the performance of the three choices (i.e., Options I, II and III in Table 2) for snapshot and starting points for solving ridge regression and Lasso problems, as shown in Figs. 3 and 4. Except for the three different settings for snapshot and starting points, we use the update rules in (12) and (13) for ridge regression and Lasso problems, respectively.

From all the results shown in Figs. 3 and 4, we can see that our algorithms with Option III (i.e., Algorithms 2 and 3 with their Option I) consistently converge much faster than SVRG with the choices of Options I and II for both strongly convex and non-strongly convex cases. This indicates that the setting of Option III suggested in this paper is a better choice than Options I and II for stochastic optimization.

### 5.3 Logistic Regression

In this part, we focus on the following generalized logistic regression problems for binary classification,

$$\min_{x\in\mathbb{R}^d} \frac{1}{n}\sum_{i=1}^{n} \log(1+\exp(-b_i a_i^T x)) + \frac{\lambda_1}{2}\|x\|^2 + \lambda_2\|x\|_1, \tag{27}$$

where $\{(a_i, b_i)\}$ is a set of training examples, and $\lambda_1, \lambda_2 \geq 0$ are the regularization parameters. Note that when $\lambda_2 > 0$, $f_i(x) = \log(1+\exp(-b_i a_i^T x)) + (\lambda_1/2)\|x\|^2$. The formulation (27) includes the $\ell_2$-norm (i.e., $\lambda_2 = 0$), $\ell_1$-norm (i.e., $\lambda_1 = 0$), and elastic net (i.e., $\lambda_1 \neq 0$ and $\lambda_2 \neq 0$) regularized logistic regression problems. Figs. 5, 6 and 7 show how the objective gap, i.e., $F(x^s) - F(x^*)$, decreases for the $\ell_2$-norm, $\ell_1$-norm, and elastic net regularized logistic regression problems, respectively. From all the results, we make the following observations.

6. As Options I and II in Algorithms 2 and 3 achieve very similar performance, we only report the results of Algorithms 2 and 3 with Option I.

- When the regularization parameters $\lambda_1$ and $\lambda_2$ are relatively large, e.g., $\lambda_1 = 10^{-4}$ or $\lambda_2 = 10^{-4}$, Prox-SVRG usually converges faster than SVRG for both strongly convex (e.g., $\ell_2$-norm regularized logistic regression) and non-strongly convex (e.g., $\ell_1$-norm regularized logistic regression) cases, as shown in Figs. 5(a)–5(c) and Figs. 6(a)–6(b). On the contrary, SVRG often outperforms Prox-SVRG, when the regularization parameters are relatively small, e.g., $\lambda_1 = 10^{-6}$ or $\lambda_2 = 10^{-6}$, as observed in [45]. The main reason is that they have different initialization settings, i.e., $\widetilde{x}^s = x_m^s$ and $x_0^{s+1} = x_m^s$ for SVRG vs. $\widetilde{x}^s = \frac{1}{m}\sum_{k=1}^m x_k^s$ and $x_0^{s+1} = \frac{1}{m}\sum_{k=1}^m x_k^s$ for Prox-SVRG.
- Katyusha converges much faster than SVRG and Prox-SVRG for the cases when the regularization parameters are relatively small, e.g., $\lambda_1 = 10^{-6}$, whereas it often achieves similar or inferior performance when the regularization parameters are relatively large, e.g., $\lambda_1 = 10^{-4}$, as shown in Figs. 5(a)–5(d) and Figs. 6(a)–6(b). Note that we implemented the original algorithms with Option I in [3] for Katyusha. In other words, Katyusha is an accelerated proximal stochastic gradient method. Obviously, the above observation matches the convergence properties of Katyusha provided in [3], that is, only if $m\mu/L \leq 3/4$, Katyusha attains the best known overall complexities of $\mathcal{O}((n+\sqrt{nL/\mu})\log(1/\epsilon))$ for strongly convex problems.
- Our VR-SGD method consistently converges much faster than SVRG and Prox-SVRG, especially when the regularization parameters are relatively small, e.g., $\lambda_1 = 10^{-6}$ or $\lambda_2 = 10^{-6}$, as shown in Figs. 5(i)–5(l) and Figs. 6(e)–6(f). The main reason is that VR-SGD can use much larger learning rates than SVRG (e.g., $6/(5L)$ for VR-SGD vs. $1/(10L)$ for SVRG), which leads to faster convergence. This further verifies that the setting of both snapshot and starting points in our algorithms (i.e., Algorithms 2 and 3) is a better choice than Options I and II in Algorithm 1.
- In particular, VR-SGD generally outperforms the best-known stochastic method, Katyusha, in terms of the number of passes through the data, especially when the regularization parameters are relatively large, e.g., $10^{-4}$ and $10^{-5}$, as shown in Figs. 5(a)–5(h) and Figs. 6(a)–6(d). Since VR-SGD has a much lower per-iteration complexity than Katyusha, VR-SGD has more obvious advantage over Katyusha in terms of running time. From the algorithms of Katyusha proposed in [3], one can see that the learning rate of Katyusha is at least set to $1/(3L)$. Similarly, the learning rate used in VR-SGD is comparable to Katyusha, which may be the main reason why the performance of VR-SGD is much better than that of Katyusha. This also implies that the algorithm that enjoys larger learning rates can yield better performance.

### 5.4 Common Stochastic Gradient and Prox-SG Updates

In this part, we compare the original Katyusha algorithm, i.e., Algorithm 1 in [3], with the slightly modified Katyusha algorithm (denoted by Katyusha-I). In Katyusha-I, only the following two update rules for smooth objective functions are used to replace the original proximal stochastic gradient update rules in (6b) and (6c).

$$\begin{aligned} y_{k+1}^s &= y_k^s - \eta[\widetilde{\nabla} f_{i_k}(x_{k+1}^s) + \nabla g(x_{k+1}^s)], \\ z_{k+1}^s &= x_{k+1}^s - [\widetilde{\nabla} f_{i_k}(x_{k+1}^s) + \nabla g(x_{k+1}^s)]/(3L). \end{aligned} \tag{28}$$

Similarly, we also implement the proximal versions[7] for the original SVRG (also called SVRG-I) and the proposed VR-SGD (denoted by VR-SGD-I) methods, and denote their proximal variants by SVRG-II and VR-SGD-II, respectively. Here, the original Katyusha method is denoted by Katyusha-II.

Figs. 8 and 9 show the performance of Katyusha-I and Katyusha-II for solving ridge regression problems on the two popular data sets: Adult and Covtype. We also report the results of SVRG, VR-SGD, and their proximal variants. It is clear that Katyusha-I usually performs better than Katyusha-II (i.e., the original proximal stochastic method, Katyusha [3]), and converges significantly faster for the case when the regularization parameter is $10^{-4}$. This seems to be the main reason why Katyusha has inferior performance when the regularization parameter is relatively large, as shown in Section 5.3. In

7. Here, the proximal variant of SVRG is different from Prox-SVRG [2], and their main difference is the choices of both the snapshot point and starting point. That is, the two vectors of the former are set to the last iterate $x_m^s$, while those of Prox-SVRG are set to the average point of the previous epoch, i.e., $\frac{1}{m}\sum_{k=1}^m x_k^s$.

contrast, VR-SGD and its proximal variant have similar performance, and the former slightly outperforms the latter in most cases, as well as SVRG vs. its proximal variant. All this suggests that stochastic gradient update rules as in (12) and (28) are better choices than proximal stochastic gradient update rules as in (10), (6b) and (6c) for smooth objective functions. We also believe that our new insight can help us to design accelerated stochastic optimization methods. Both Katyusha-I and Katyusha-II usually outperform SVRG and its proximal variant, especially when the regularization parameter is relatively small, e.g., $\lambda = 10^{-6}$, as shown in Figs. 8(d) and 9(d). Unfortunately, both Katyusha-I and Katyusha-II cannot solve the convex objectives without any regularization term (i.e., $\lambda = 0$, as shown in Figs. 8(f) and 9(f)) or with too large and small regularization parameters, e.g., $10^{-3}$ and $10^{-7}$, as shown in Figs. 9(a) and 8(e).

Moreover, it can be seen that both VR-SGD and its proximal variant achieve much better performance than the other methods in most cases, and are also comparable to Katyusha-I and Katyusha-II in the remaining cases. This further verifies that VR-SGD is suitable for various large-scale machine learning.

### 5.5 Fixed and Varied Learning Rates

Finally, we compare the performance of Algorithm 3 with fixed and varied learning rates for solving $\ell_1$-norm regularized logistic regression and Lasso problems, as shown in Fig. 10. Note that the learning rates in Algorithm 3 are varied according to the update formula in (14). We can observe that Algorithm 3 with varied step-sizes performs very similar to Algorithm 3 with fixed step-sizes in most cases. In the remaining cases, the former slightly outperforms the latter.

## 6 Conclusions

In this paper, we proposed a simple variant of the original SVRG [1], called variance reduced stochastic gradient descent (VR-SGD). Unlike the choices of the snapshot point and starting point in SVRG and its proximal variant, Prox-SVRG [2], the two points of each epoch in VR-SGD are set to the average and last iterate of the previous epoch, respectively. This setting allows us to use much larger learning rates than SVRG, e.g., $3/(7L)$ for VR-SGD vs. $1/(10L)$ for SVRG, and also makes VR-SGD more robust in terms of learning rate selection. Different from existing proximal stochastic methods such as Prox-SVRG [2] and Katyusha [3], we designed two different update rules for smooth and non-smooth objective functions, respectively, which makes our VR-SGD method suitable for non-smooth and/or non-strongly convex problems without using any reduction techniques as in [3], [61]. In contrast, SVRG and Prox-SVRG cannot directly solve non-strongly convex objectives [42]. Furthermore, our empirical results showed that for smooth problems stochastic gradient update rules as in (12) are better choices than proximal stochastic gradient update formulas as in (10).

On the practical side, the choices of the snapshot and starting points make VR-SGD significantly faster than its counterparts, SVRG and Prox-SVRG. On the theoretical side, the setting also makes our convergence analysis more challenging. We analyzed the convergence properties of VR-SGD for strongly convex objective functions, which show that VR-SGD attains a linear convergence rate. Moreover, we also provided the convergence guarantees of VR-SGD for non-strongly convex problems, which show that VR-SGD achieves a sub-linear convergence rate. All these results imply that VR-SGD is guaranteed to have a similar convergence rate with other variance reduced stochastic methods such as SVRG and Prox-SVRG, and a slower theoretical rate than accelerated methods such as Katyusha [3]. Nevertheless, various experimental results show that VR-SGD significantly outperforms than SVRG and Prox-SVRG, and is still much better than the best known stochastic method, Katyusha [3].

## Appendix A: Proof of Lemma 2

Before proving Lemma 2, we first give and prove the following lemma.

**Lemma 9.** *Suppose each $f_i(\cdot)$ is $L$-smooth, and let $x^*$ be the optimal solution of Problem (1) when $g(x) \equiv 0$, then we have*

$$\mathbb{E}\left[\|\nabla f_i(x) - \nabla f_i(x^*)\|^2\right] \leq 2L\left[f(x) - f(x^*)\right].$$

*Proof.* Following Theorem 2.1.5 in [19] and Lemma 3.4 [2], we have

$$\|\nabla f_i(x) - \nabla f_i(x^*)\|^2 \leq 2L\left[f_i(x) - f_i(x^*) - \langle \nabla f_i(x^*),\ x - x^*\rangle\right].$$

Summing the above inequality over $i = 1, \ldots, n$, we obtain

$$\mathbb{E}\Big[\|\nabla f_i(x) - \nabla f_i(x^*)\|^2\Big] = \frac{1}{n}\sum_{i=1}^{n}\|\nabla f_i(x) - \nabla f_i(x^*)\|^2 \leq 2L\left[f(x) - f(x^*) - \langle \nabla f(x^*),\ x - x^*\rangle\right].$$

By the optimality of $x^*$, i.e., $x^* = \arg\min_x f(x)$, we have $\nabla f(x^*) = 0$. Then

$$\begin{aligned}\mathbb{E}\Big[\|\nabla f_i(x) - \nabla f_i(x^*)\|^2\Big] \leq & 2L\left[f(x) - f(x^*) - \langle \nabla f(x^*), x - x^*\rangle\right] \\ = & 2L\left[f(x) - f(x^*)\right].\end{aligned}$$

□

**Proof of Lemma 2:**

*Proof.*

$$\begin{aligned}\mathbb{E}\Big[\|\widetilde{\nabla} f_{i_k^s}(x_k^s) - \nabla f(x_k^s)\|^2\Big] = & \mathbb{E}\left[\left\|\nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1}) + \nabla f(\widetilde{x}^{s-1}) - \nabla f(x_k^s)\right\|^2\right] \\ = & \mathbb{E}\left[\left\|\nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1})\right\|^2\right] - \left\|\nabla f(x_k^s) - \nabla f(\widetilde{x}^{s-1})\right\|^2 \\ \leq & \mathbb{E}\left[\left\|\nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1})\right\|^2\right] \\ \leq & 2\mathbb{E}\left[\left\|\nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(x^*)\right\|^2\right] + 2\mathbb{E}\left[\left\|\nabla f_{i_k^s}(\widetilde{x}^{s-1}) - \nabla f_{i_k^s}(x^*)\right\|^2\right] \\ \leq & 4L\big[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\big],\end{aligned}$$

where the second equality holds due to the fact that $\mathbb{E}[\|x - \mathbb{E}[x]\|^2] = \mathbb{E}[\|x\|^2] - \|\mathbb{E}[x]\|^2$; the second inequality holds due to the fact that $\|a - b\|^2 \leq 2(\|a\|^2 + \|b\|^2)$; and the last inequality follows from Lemma 9. □

## APPENDIX B: PROOF OF LEMMA 5

*Proof.* For convenience, the stochastic gradient estimator is defined as: $v_k^s := \nabla f_{i_k^s}(x_k^s) - \nabla f_{i_k^s}(\widetilde{x}^{s-1}) + \nabla f(\widetilde{x}^{s-1})$. Since each component function $f_i(x)$ is $L$-smooth, which implies that the gradient of the average function $f(x)$ is also $L$-smooth, i.e., for all $x, y \in \mathbb{R}^d$,

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|,$$

whose equivalent form is

$$f(y) \leq f(x) + \langle \nabla f(x),\ y - x\rangle + \frac{L}{2}\|y - x\|^2.$$

Using the above smoothness inequality, we have

$$\begin{aligned}f(x_{k+1}^s) \leq & f(x_k^s) + \langle \nabla f(x_k^s),\ x_{k+1}^s - x_k^s\rangle + \frac{L}{2}\|x_{k+1}^s - x_k^s\|^2 \\ = & f(x_k^s) + \langle \nabla f(x_k^s),\ x_{k+1}^s - x_k^s\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2 - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2 \\ = & f(x_k^s) + \langle v_k^s,\ x_{k+1}^s - x_k^s\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2 \\ & + \langle \nabla f(x_k^s) - v_k^s,\ x_{k+1}^s - x_k^s\rangle - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2.\end{aligned} \tag{29}$$

Using Lemma 2, then we get

$$\begin{aligned}& \mathbb{E}\left[\langle \nabla f(x_k^s) - v_k^s,\ x_{k+1}^s - x_k^s\rangle - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2\right] \\ \leq & \mathbb{E}\left[\frac{1}{2L(\beta-1)}\|\nabla f(x_k^s) - v_k^s\|^2 + \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2 - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2\right] \\ \leq & \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right],\end{aligned} \tag{30}$$

where the first inequality holds due to the Young's inequality (i.e., $y^T z \le \|y\|^2/(2\theta)+\theta\|z\|^2/2$ for all $\theta>0$ and $y, z\in\mathbb{R}^d$), and the second inequality follows from Lemma 2.

Taking the expectation over the random choice of $i_k^s$ and substituting the inequality in (30) into the inequality in (29), we have

$$\begin{aligned}
&\mathbb{E}[f(x_{k+1}^s)]\\
&\le \mathbb{E}[f(x_k^s)] + \mathbb{E}\left[\langle v_k^s,\, x_{k+1}^s - x_k^s\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2\right] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right]\\
&\le \mathbb{E}[f(x_k^s)] + \mathbb{E}\left[\langle v_k^s,\, x^* - x_k^s\rangle + \frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right]\\
&\le \mathbb{E}[f(x_k^s)] + \langle \nabla f(x_k^s),\, x^* - x_k^s\rangle + \mathbb{E}\left[\frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right]\\
&\le f(x^*) + \mathbb{E}\left[\frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right]\\
&= f(x^*) + \frac{L\beta}{2}\mathbb{E}\left[(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] + \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right],
\end{aligned}$$

where the first inequality holds due to the inequality in (29) and the inequality in (30); the second inequality follows from Lemma 1 with $\hat{z} = x_{k+1}^s$, $z = x^*$, $z_0 = x_k^s$, $\tau = L\beta = 1/\eta$, and $r(z) := \langle v_k^s,\, z - x_k^s\rangle$; the third inequality holds due to the fact that $\mathbb{E}[v_k^s] = \nabla f(x_k^s)$; and the last inequality follows from the convexity of $f(\cdot)$, i.e., $f(x_k^s) + \langle \nabla f(x_k^s),\, x^* - x_k^s\rangle \le f(x^*)$. The above inequality can be rewritten as follows:

$$\mathbb{E}[f(x_{k+1}^s)] - f(x^*) \le \frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right] + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2\right].$$

Summing the above inequality over $k = 0, 1, \ldots, m-1$, we obtain

$$\begin{aligned}
&\sum_{k=0}^{m-1}\left\{\mathbb{E}[f(x_{k+1}^s)] - f(x^*)\right\}\\
&\le \sum_{k=0}^{m-1}\left\{\frac{2}{\beta-1}\left[f(x_k^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right] + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2\right]\right\}.
\end{aligned}$$

Then

$$\begin{aligned}
&\left(1 - \frac{2}{\beta-1}\right)\sum_{k=1}^{m}\left\{\mathbb{E}[f(x_k^s)] - f(x^*)\right\} + \frac{2}{\beta-1}\sum_{k=1}^{m}\left\{\mathbb{E}[f(x_k^s)] - f(x^*)\right\}\\
&\le \sum_{k=1}^{m}\left\{\frac{2}{\beta-1}\left[f(x_{k-1}^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right]\right\} + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned} \tag{31}$$

Due to the setting of $\widetilde{x}^s = \frac{1}{m}\sum_{k=1}^m x_k^s$ in Option II, and the convexity of $f(\cdot)$, then

$$f(\widetilde{x}^s) \le \frac{1}{m}\sum_{k=1}^{m} f(x_k^s).$$

Using the above inequality, the inequality in (31) becomes

$$\begin{aligned}
&m\left(1 - \frac{2}{\beta-1}\right)\mathbb{E}[f(\widetilde{x}^s) - f(x^*)] + \frac{2}{\beta-1}\sum_{k=1}^{m}\left\{\mathbb{E}[f(x_k^s)] - f(x^*)\right\}\\
&\le \sum_{k=1}^{m}\left\{\frac{2}{\beta-1}\left[f(x_{k-1}^s) - f(x^*) + f(\widetilde{x}^{s-1}) - f(x^*)\right]\right\} + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

Dividing both sides of the above inequality by $m$ and subtracting $\frac{2}{(\beta-1)m}\sum_{k=1}^{m-1}[f(x_k^s) - f(x^*)]$ from both sides, we arrive at

$$\begin{aligned}
&\left(1 - \frac{2}{\beta-1}\right)\mathbb{E}[f(\widetilde{x}^s) - f(x^*)] + \frac{2}{(\beta-1)m}\mathbb{E}[f(x_m^s) - f(x^*)]\\
&\le \frac{2}{(\beta-1)}\mathbb{E}\left[f(\widetilde{x}^{s-1}) - f(x^*)\right] + \frac{2}{(\beta-1)m}\mathbb{E}[f(x_0^s) - f(x^*)] + \frac{L\beta}{2m}\mathbb{E}\left[\|x^* - x_0^s\|^2 - \|x^* - x_m^s\|^2\right].
\end{aligned}$$

This completes the proof. □

## APPENDIX C: PROOF OF THEOREM 2

*Proof.* Using Lemma 5, we have

$$\begin{aligned}&(1-\frac{2}{\beta-1})\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]+\frac{2}{(\beta-1)m}\mathbb{E}[f(x_m^s)-f(x^*)]\\ \le&\frac{2}{(\beta-1)m}\mathbb{E}[f(x_0^s)-f(x^*)]+\frac{2}{\beta-1}\mathbb{E}\big[f(\widetilde{x}^{s-1})-f(x^*)\big]+\frac{L\beta}{2m}\mathbb{E}[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2].\end{aligned}$$

Summing the above inequality over $s=1,2,\ldots,S$, taking expectation with respect to the history of random variables $i_k^s$, and using the setting of $x_0^{s+1}=x_m^s$ in Algorithm 3, we arrive at

$$\begin{aligned}&\sum_{s=1}^{S}(1-\frac{2}{\beta-1})\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]\\ \le&\sum_{s=1}^{S}\left\{\frac{2}{(\beta-1)m}\mathbb{E}\{f(x_0^s)-f(x^*)-[f(x_m^s)-f(x^*)]\}+\frac{2}{\beta-1}\mathbb{E}\big[f(\widetilde{x}^{s-1})-f(x^*)\big]\right\}\\ &+\frac{L\beta}{2m}\sum_{s=1}^{S}\mathbb{E}\big[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2\big].\end{aligned}$$

Subtracting $\frac{2}{\beta-1}\sum_{s=1}^{S}[f(\widetilde{x}^s)-f(x^*)]$ from both sides of the above inequality, we obtain

$$\begin{aligned}&\sum_{s=1}^{S}\left(1-\frac{4}{\beta-1}\right)\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]\\ \le&\frac{2}{(\beta-1)m}\mathbb{E}\Big\{f(x_0^1)-f(x^*)-[f(x_m^S)-f(x^*)]\Big\}+\frac{2}{\beta-1}\mathbb{E}[f(\widetilde{x}^0)-f(\widetilde{x}^S)]\\ &+\frac{L\beta}{2m}\mathbb{E}\Big[\|x^*-x_0^1\|^2-\|x^*-x_m^S\|^2\Big].\end{aligned}$$

It is not hard to verify that $\mathbb{E}[f(\widetilde{x}^0)-f(\widetilde{x}^S)]\le f(\widetilde{x}^0)-f(x^*)$. Dividing both sides of the above inequality by $S$, and using the choice $\widetilde{x}^0=x_0^1$, we have

$$\begin{aligned}&\frac{1}{S}\left(1-\frac{4}{\beta-1}\right)\sum_{s=1}^{S}\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]\\ \le&\frac{2}{(\beta-1)mS}\mathbb{E}\Big\{f(x_0^1)-f(x^*)-[f(x_m^S)-f(x^*)]\Big\}+\frac{2}{(\beta-1)S}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta}{2mS}\|x^*-x_0^1\|^2\\ \le&\frac{2}{(\beta-1)mS}\big[f(\widetilde{x}^0)-f(x^*)\big]+\frac{2}{(\beta-1)S}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta}{2mS}\|x^*-\widetilde{x}^0\|^2\\ =&\frac{2(m+1)}{(\beta-1)mS}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta}{2mS}\|\widetilde{x}^0-x^*\|^2\end{aligned}\tag{32}$$

where the first inequality holds due to the facts that $\mathbb{E}[f(\widetilde{x}^0)-f(\widetilde{x}^S)]\le f(\widetilde{x}^0)-f(x^*)$ and $\mathbb{E}\big[\|x^*-x_m^S\|^2\big]\ge 0$, and the last inequality uses the facts that $\mathbb{E}\big[f(x_m^S)-f(x^*)\big]\ge 0$ and $\widetilde{x}^0=x_0^1$.

Since $\overline{x}^S=\frac{1}{S}\sum_{s=1}^{S}\widetilde{x}^s$, and using the convexity of $f(\cdot)$, we have $f(\overline{x}^S)\le\frac{1}{S}\sum_{s=1}^{S}f(\widetilde{x}^s)$, and therefore the inequality in (32) becomes

$$\begin{aligned}\left(1-\frac{4}{\beta-1}\right)\mathbb{E}\Big[f(\overline{x}^S)-f(x^*)\Big]&\le\frac{1}{S}\left(1-\frac{4}{\beta-1}\right)\sum_{s=1}^{S}\mathbb{E}[f(\widetilde{x}^s)-f(x^*)]\\ &\le\frac{2(m+1)}{(\beta-1)mS}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta}{2mS}\|\widetilde{x}^0-x^*\|^2.\end{aligned}$$

Dividing both sides of the above inequality by $(1-\frac{4}{\beta-1})>0$ (i.e., $\eta<1/(5L)$), we arrive at

$$\mathbb{E}\Big[f(\overline{x}^S)-f(x^*)\Big]\le\frac{2(m+1)}{mS(\beta-5)}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta(\beta-1)}{2mS(\beta-5)}\|\widetilde{x}^0-x^*\|^2.$$

If $f(\widetilde{x}^S)\le f(\overline{x}^S)$, then $\widehat{x}^S=\widetilde{x}^S$, and

$$\begin{aligned}\mathbb{E}\Big[f(\widehat{x}^S)-f(x^*)\Big]&\le\mathbb{E}\Big[f(\overline{x}^S)-f(x^*)\Big]\\ &\le\frac{2(m+1)}{mS(\beta-5)}[f(\widetilde{x}^0)-f(x^*)]+\frac{L\beta(\beta-1)}{2mS(\beta-5)}\|\widetilde{x}^0-x^*\|^2.\end{aligned}$$

Alternatively, if $f(\widetilde{x}^S) \geq f(\overline{x}^S)$, then $\widehat{x}^S = \overline{x}^S$, and the above inequality still holds.

This completes the proof. □

## APPENDIX D: PROOF OF LEMMA 6

*Proof.* Since the average function $f(x)$ is $L$-smooth, then for all $x, y \in \mathbb{R}^d$,

$$f(y) \leq f(x) + \langle \nabla f(x),\ y - x \rangle + \frac{L}{2}\|y - x\|^2,$$

which then implies

$$f(x_{k+1}^s) \leq f(x_k^s) + \left\langle \nabla f(x_k^s),\ x_{k+1}^s - x_k^s \right\rangle + \frac{L}{2}\left\|x_{k+1}^s - x_k^s\right\|^2.$$

Using the above inequality, we have

$$\begin{aligned} F(x_{k+1}^s) &= f(x_{k+1}^s) + g(x_{k+1}^s) \\ &\leq f(x_k^s) + g(x_{k+1}^s) + \left\langle \nabla f(x_k^s),\ x_{k+1}^s - x_k^s \right\rangle + \frac{L\beta}{2}\left\|x_{k+1}^s - x_k^s\right\|^2 - \frac{L(\beta-1)}{2}\left\|x_{k+1}^s - x_k^s\right\|^2 \\ &= f(x_k^s) + g(x_{k+1}^s) + \left\langle v_k^s,\ x_{k+1}^s - x_k^s \right\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2 \\ &\quad + \left\langle \nabla f(x_k^s) - v_k^s,\ x_{k+1}^s - x_k^s \right\rangle - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2. \end{aligned} \tag{33}$$

According to Lemma 5, then we obtain

$$\begin{aligned} &\mathbb{E}\left[\left\langle \nabla f(x_k^s) - v_k^s,\ x_{k+1}^s - x_k^s \right\rangle - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2\right] \\ \leq &\mathbb{E}\left[\frac{1}{2L(\beta-1)}\|\nabla f(x_k^s) - v_k^s\|^2 + \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2 - \frac{L(\beta-1)}{2}\|x_{k+1}^s - x_k^s\|^2\right] \\ \leq &\frac{2}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right], \end{aligned} \tag{34}$$

where the first inequality holds due to the Young's inequality, and the second inequality follows from Lemma 5. Substituting the inequality (34) into the inequality (33), and taking the expectation over the random choice $i_k^s$, we arrive at

$$\begin{aligned} \mathbb{E}\left[F(x_{k+1}^s)\right] &\leq \mathbb{E}[f(x_k^s)] + \mathbb{E}\left[g(x_{k+1}^s)\right] + \mathbb{E}\left[\left\langle v_k^s,\ x_{k+1}^s - x_k^s \right\rangle + \frac{L\beta}{2}\|x_{k+1}^s - x_k^s\|^2\right] \\ &\quad + \frac{2}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] \\ &\leq \mathbb{E}[f(x_k^s)] + g(x^*) + \mathbb{E}\left[\left\langle v_k^s,\ x^* - x_k^s \right\rangle + \frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] \\ &\quad + \frac{2}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] \\ &\leq f(x^*) + g(x^*) + \mathbb{E}\left[\frac{L\beta}{2}(\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2)\right] \\ &\quad + \frac{2}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] \\ &= F(x^*) + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2\right] + \frac{2}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right], \end{aligned}$$

where the first inequality holds due to the inequality (33) and the inequality (34); the second inequality follows from Lemma 1 with $\hat{z} = x_{k+1}^s$, $z = x^*$, $z_0 = x_k^s$, $\tau = L\beta = 1/\eta$, and $r(z) := \langle v_k^s,\ z - x_k^s \rangle + g(z)$; and the third inequality holds due to the fact that $\mathbb{E}[v_k^s] = \nabla f(x_k^s)$ and the convexity of $f(\cdot)$, i.e., $f(x_k^s) + \langle \nabla f(x_k^s),\ x^* - x_k^s \rangle \leq f(x^*)$. Then the above inequality is rewritten as follows:

$$\begin{aligned} &\mathbb{E}[F(x_{k+1}^s)] - F(x^*) \\ \leq &\frac{2}{\beta-1}\left[F(x_k^s) - F(x^*) + F(\widetilde{x}^{s-1}) - F(x^*)\right] + \frac{L\beta}{2}\mathbb{E}\left[\|x^* - x_k^s\|^2 - \|x^* - x_{k+1}^s\|^2\right]. \end{aligned} \tag{35}$$

Summing the above inequality over $k=0,1,\ldots,(m-1)$ and taking expectation over whole history, we have

$$\begin{aligned}&\sum_{k=0}^{m-1}\left\{\mathbb{E}[F(x_{k+1}^s)]-F(x^*)\right\}\\ \leq&\sum_{k=0}^{m-1}\left\{\frac{2}{\beta-1}\left[F(x_k^s)-F(x^*)+F(\widetilde{x}^{s-1})-F(x^*)\right]+\frac{L\beta}{2}\mathbb{E}\left[\|x^*-x_k^s\|^2-\|x^*-x_{k+1}^s\|^2\right]\right\}.\end{aligned}$$

Subtracting $\frac{2}{\beta-1}\sum_{k=0}^{m-2}\mathbb{E}\left[F(x_{k+1}^s)-F(x^*)\right]$ from both sides of the above inequality, we obtain

$$\begin{aligned}&\sum_{k=0}^{m-1}\mathbb{E}\left[F(x_{k+1}^s)-F(x^*)\right]-\frac{2}{\beta-1}\sum_{k=0}^{m-1}\mathbb{E}\left[F(x_{k+1}^s)-F(x^*)\right]+\frac{2}{\beta-1}\mathbb{E}[F(x_m^s)-F(x^*)]\\ \leq&\sum_{k=0}^{m-1}\left\{\frac{2}{\beta-1}\left[F(x_k^s)-F(x^*)+F(\widetilde{x}^{s-1})-F(x^*)\right]\right\}-\frac{2}{\beta-1}\sum_{k=0}^{m-2}\mathbb{E}\left[F(x_{k+1}^s)-F(x^*)\right]+\frac{L\beta}{2}\mathbb{E}\left[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2\right].\end{aligned}$$

Then

$$\begin{aligned}&\left(1-\frac{2}{\beta-1}\right)\sum_{k=1}^{m}\mathbb{E}[F(x_k^s)-F(x^*)]+\frac{2}{\beta-1}\mathbb{E}[F(x_m^s)-F(x^*)]\\ \leq&\frac{2}{\beta-1}\mathbb{E}[F(x_0^s)-F(x^*)]+\frac{2m}{\beta-1}\mathbb{E}\left[F(\widetilde{x}^{s-1})-F(x^*)\right]+\frac{L\beta}{2}\mathbb{E}\left[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2\right].\end{aligned}$$

Due to the settings of $\widetilde{x}^s=\frac{1}{m}\sum_{k=1}^m x_k^s$ and $x_0^{s+1}=x_m^s$, and the convexity of the objective function $F(\cdot)$, we have $F(\widetilde{x}^s)\leq\frac{1}{m}\sum_{k=1}^m F(x_k^s)$, and

$$\begin{aligned}&m\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]+\frac{2}{\beta-1}\mathbb{E}[F(x_m^s)-F(x^*)]\\ \leq&\frac{2}{\beta-1}\mathbb{E}[F(x_0^s)-F(x^*)]+\frac{2m}{\beta-1}\mathbb{E}\left[F(\widetilde{x}^{s-1})-F(x^*)\right]+\frac{L\beta}{2}\mathbb{E}\left[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2\right].\end{aligned}$$

Dividing both sides of the above inequality by $m$, we arrive at

$$\begin{aligned}&\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]+\frac{2}{(\beta-1)m}\mathbb{E}[F(x_m^s)-F(x^*)]\\ \leq&\frac{2}{(\beta-1)m}\mathbb{E}[F(x_0^s)-F(x^*)]+\frac{2}{\beta-1}\mathbb{E}\left[F(\widetilde{x}^{s-1})-F(x^*)\right]+\frac{L\beta}{2m}\mathbb{E}\left[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2\right].\end{aligned}$$

This completes the proof. □

## APPENDIX E: PROOF OF THEOREM 3

*Proof.* Summing the inequality in (21) over $s=1,2,\ldots,S$, and taking expectation with respect to the history of $i_k^s$, we have

$$\begin{aligned}&\sum_{s=1}^{S}\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]+\sum_{s=1}^{S}\frac{2}{(\beta-1)m}\mathbb{E}[F(x_m^s)-F(x^*)]\\ \leq&\sum_{s=1}^{S}\left\{\frac{2}{(\beta-1)m}\mathbb{E}[F(x_0^s)-F(x^*)]+\frac{2}{\beta-1}\mathbb{E}\left[F(\widetilde{x}^{s-1})-F(x^*)\right]\right\}\\ &+\frac{L\beta}{2m}\sum_{s=1}^{S}\mathbb{E}\left[\|x^*-x_0^s\|^2-\|x^*-x_m^s\|^2\right].\end{aligned}$$

Subtracting $\sum_{s=1}^S\frac{2}{(\beta-1)m}\mathbb{E}[F(x_m^s)-F(x^*)]+\frac{2}{\beta-1}\sum_{s=1}^S[F(\widetilde{x}^s)-F(x^*)]$ from both sides of the above inequality, and using the setting of $x_0^{s+1}=x_m^s$, we arrive at

$$\begin{aligned}&\sum_{s=1}^{S}\left(1-\frac{4}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]\\ \leq&\frac{2}{(\beta-1)m}\mathbb{E}\Big[F(x_0^1)-F(x_m^S)\Big]+\frac{2}{\beta-1}\mathbb{E}\Big[F(\widetilde{x}^0)-F(\widetilde{x}^S)\Big]+\frac{L\beta}{2m}\mathbb{E}\Big[\|x^*-x_0^1\|^2-\|x^*-x_m^S\|^2\Big].\end{aligned}$$

It is easy to verify that $\mathbb{E}[F(\widetilde{x}^0)-F(\widetilde{x}^S)]\leq F(\widetilde{x}^0)-F(x^*)$. Dividing both sides of the above inequality by $S$, and using the choice $\widetilde{x}^0=x_0^1$, we obtain

$$\begin{aligned}
&\left(1-\frac{4}{\beta-1}\right)\frac{1}{S}\sum_{s=1}^{S}\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]\\
&\leq\frac{2}{(\beta-1)mS}\mathbb{E}\Big[F(x_0^1)-F(x_m^S)\Big]+\frac{2}{(\beta-1)S}[F(\widetilde{x}^0)-F(\widetilde{x}^S)]+\frac{L\beta}{2mS}\|x^*-x_0^1\|^2\\
&\leq\frac{2}{(\beta-1)mS}\big[F(\widetilde{x}^0)-F(x^*)\big]+\frac{2}{(\beta-1)S}[F(\widetilde{x}^0)-F(x^*)]+\frac{L\beta}{2mS}\|x^*-\widetilde{x}^0\|^2\\
&=\frac{2(m+1)}{(\beta-1)mS}[F(\widetilde{x}^0)-F(x^*)]+\frac{L\beta}{2mS}\|\widetilde{x}^0-x^*\|^2
\end{aligned}\tag{36}$$

where the first inequality uses the fact that $\|x^*-x_m^S\|^2\geq 0$; and the last inequality holds due to the facts that $\mathbb{E}\big[F(x_0^1)-F(x_m^S)\big]\leq F(x_0^1)-F(x^*)$, $\mathbb{E}[F(\widetilde{x}^0)-F(\widetilde{x}^S)]\leq F(\widetilde{x}^0)-F(x^*)$, and $\widetilde{x}^0=x_0^1$.

Using the definition of $\overline{x}^S=\frac{1}{S}\sum_{s=1}^{S}\widetilde{x}^s$ and the convexity of the objective function $F(\cdot)$, we have $F(\overline{x}^S)\leq\frac{1}{S}\sum_{s=1}^{S}F(\widetilde{x}^s)$, and therefore we can rewrite the above inequality in (36) as

$$\begin{aligned}
\left(1-\frac{4}{\beta-1}\right)\mathbb{E}\Big[F(\overline{x}^S)-F(x^*)\Big]&\leq\left(1-\frac{4}{\beta-1}\right)\frac{1}{S}\sum_{s=1}^{S}\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]\\
&\leq\frac{2(m+1)}{(\beta-1)mS}[F(\widetilde{x}^0)-F(x^*)]+\frac{L\beta}{2mS}\|\widetilde{x}^0-x^*\|^2.
\end{aligned}$$

Dividing both sides of the above inequality by $(1-\frac{4}{\beta-1})>0$, we have

$$\mathbb{E}\Big[F(\overline{x}^S)\Big]-F(x^*)\leq\frac{2(m+1)}{(\beta-5)mS}[F(\widetilde{x}^0)-F(x^*)]+\frac{\beta(\beta-1)L}{2(\beta-5)mS}\|\widetilde{x}^0-x^*\|^2.$$

When $F(\widetilde{x}^S)\leq F(\overline{x}^S)$, then $\widehat{x}^S=\widetilde{x}^S$, and

$$\mathbb{E}\Big[F(\widehat{x}^S)\Big]-F(x^*)\leq\frac{2(m+1)}{(\beta-5)mS}[F(\widetilde{x}^0)-F(x^*)]+\frac{\beta(\beta-1)L}{2(\beta-5)mS}\|\widetilde{x}^0-x^*\|^2.$$

Alternatively, if $F(\widetilde{x}^S)\geq F(\overline{x}^S)$, then $\widehat{x}^S=\overline{x}^S$, and the above inequality still holds.

This completes the proof. □

## APPENDIX F: PROOF OF THEOREM 6

*Proof.* Since each $f_i(\cdot)$ is convex and $L$-smooth, then we have

$$\begin{aligned}
&\left(1-\frac{2}{\beta-1}\right)\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]+\frac{2}{m(\beta-1)}\mathbb{E}\big[F(x_0^{s+1})-F(x^*)\big]+\frac{L\beta}{2m}\mathbb{E}\big[\|x_0^{s+1}-x^*\|^2\big]\\
&\leq\frac{2}{\beta-1}\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]+\frac{2}{m(\beta-1)}\mathbb{E}[F(x_0^s)-F(x^*)]+\frac{L\beta}{2m}\mathbb{E}\big[\|x_0^s-x^*\|^2\big]\\
&\leq\frac{2}{\beta-1}\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]+\left(\frac{2}{m(\beta-1)}+\frac{L\beta}{m\mu}\right)\mathbb{E}[F(x_0^s)-F(x^*)]\\
&\leq\left(\frac{2(m+C)}{m(\beta-1)}+\frac{CL\beta}{m\mu}\right)\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big]
\end{aligned}$$

where the first inequality follows from Lemma 6; the second inequality holds due to the fact that $\|x_0^s-x^*\|^2\leq(2/\mu)[F(x_0^s)-F(x^*)]$; and the last inequality follows from Assumption 3.

Due to the definition of $\beta=1/(L\eta)$, the above inequality is rewritten as follows:

$$\frac{1-3L\eta}{1-L\eta}\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]\leq\left(\frac{2L\eta(m+C)}{m(1-L\eta)}+\frac{C}{m\mu\eta}\right)\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big].$$

Dividing both sides of the above inequality by $(1-3L\eta)(1-L\eta)>0$, we arrive at

$$\mathbb{E}[F(\widetilde{x}^s)-F(x^*)]\leq\left(\frac{2L\eta(m+C)}{m(1-3L\eta)}+\frac{C(1-L\eta)}{m\mu\eta(1-3L\eta)}\right)\mathbb{E}\big[F(\widetilde{x}^{s-1})-F(x^*)\big].$$

This completes the proof. □

## REFERENCES


[1] R. Johnson and T. Zhang, "Accelerating stochastic gradient descent using predictive variance reduction," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 315–323.

[2] L. Xiao and T. Zhang, "A proximal stochastic gradient method with progressive variance reduction," *SIAM J. Optim.*, vol. 24, no. 4, pp. 2057–2075, 2014.

[3] Z. Allen-Zhu, "Katyusha: The first direct acceleration of stochastic gradient methods," in *Proc. 49th ACM Symp. Theory Comput.*, 2017.

[4] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.

[5] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *Proc. 30th Int. Conf. Mach. Learn.*, 2013, pp. 1139–1147.

[6] Z. Allen-Zhu and E. Hazan, "Variance reduction for faster non-convex optimization," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 699–707.

[7] H. Ouyang, N. He, L. Q. Tran, and A. Gray, "Stochastic alternating direction method of multipliers," in *Proc. 30th Int. Conf. Mach. Learn.*, 2013, pp. 80–88.

[8] Y. Liu, F. Shang, and J. Cheng, "Accelerated variance reduced stochastic ADMM," in *Proc. 31st AAAI Conf. Artif. Intell.*, 2017, pp. 2287–2293.

[9] C. Qu, Y. Li, and H. Xu, "Linear convergence of SVRG in statistical estimation," *arXiv:1611.01957v2*, 2017.

[10] C. Paquette, H. Lin, D. Drusvyatskiy, J. Mairal, and Z. Harchaoui, "Catalyst acceleration for gradient-based non-convex optimization," *arXiv:1703.10993*, 2017.

[11] J. Duchi and F. Ruan, "Stochastic methods for composite optimization problems," *arXiv:1703.08570*, 2017.

[12] B. Recht and C. Ré, "Parallel stochastic gradient algorithms for large-scale matrix completion," *Math. Prog. Comp.*, vol. 5, pp. 201–226, 2013.

[13] X. Zhang, L. Wang, and Q. Gu, "Stochastic variance-reduced gradient descent for low-rank matrix recovery from linear measurements," *arXiv:1701.00481v2*, 2017.

[14] M. Schmidt, R. Babanezhad, M. Ahmed, A. Defazio, A. Clifton, and A. Sarkar, "Non-uniform stochastic average gradient method for training conditional random fields," in *Proc. 18th Int. Conf. Artif. Intell. Statist.*, 2015, pp. 819–828.

[15] O. Shamir, "A stochastic PCA and SVD algorithm with an exponential convergence rate," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 144–152.

[16] D. Garber, E. Hazan, C. Jin, S. M. Kakade, C. Musco, P. Netrapalli, and A. Sidford, "Faster eigenvector computation via shift-and-invert preconditioning," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 2626–2634.

[17] Z. Allen-Zhu and Y. Li, "Doubly accelerated methods for faster CCA and generalized eigendecomposition," *arXiv:1607.06017v2*, 2017.

[18] Y. Nesterov, "A method of solving a convex programming problem with convergence rate $O(1/k^2)$," *Soviet Math. Doklady*, vol. 27, pp. 372–376, 1983.

[19] ——, *Introductory Lectures on Convex Optimization: A Basic Course*. Boston: Kluwer Academic Publ., 2004.

[20] P. Tseng, "On accelerated proximal gradient methods for convex-concave optimization," *Technical report, University of Washington*, 2008.

[21] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, 2009.

[22] H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Statist.*, vol. 22, no. 3, pp. 400–407, 1951.

[23] T. Zhang, "Solving large scale linear prediction problems using stochastic gradient descent algorithms," in *Proc. 21st Int. Conf. Mach. Learn.*, 2004, pp. 919–926.

[24] C. Hu, J. T. Kwok, and W. Pan, "Accelerated gradient methods for stochastic optimization and online learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2009, pp. 781–789.

[25] S. Bubeck, "Convex optimization: Algorithms and complexity," *Found. Trends Mach. Learn.*, vol. 8, pp. 231–358, 2015.

[26] A. Rakhlin, O. Shamir, and K. Sridharan, "Making gradient descent optimal for strongly convex stochastic optimization," in *Proc. 29th Int. Conf. Mach. Learn.*, 2012, pp. 449–456.

[27] O. Shamir and T. Zhang, "Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes," in *Proc. 30th Int. Conf. Mach. Learn.*, 2013, pp. 71–79.

[28] N. L. Roux, M. Schmidt, and F. Bach, "A stochastic gradient method with an exponential convergence rate for finite training sets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 2672–2680.

[29] S. Shalev-Shwartz and T. Zhang, "Stochastic dual coordinate ascent methods for regularized loss minimization," *J. Mach. Learn. Res.*, vol. 14, pp. 567–599, 2013.

[30] A. Defazio, F. Bach, and S. Lacoste-Julien, "SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 1646–1654.

[31] Y. Zhang and L. Xiao, "Stochastic primal-dual coordinate method for regularized empirical risk minimization," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 353–361.

[32] M. Schmidt, N. L. Roux, and F. Bach, "Minimizing finite sums with the stochastic average gradient," INRIA, Paris, Tech. Rep., 2013.

[33] S. Shalev-Shwartz and T. Zhang, "Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization," *Math. Program.*, vol. 155, pp. 105–145, 2016.

[34] F. Shang, Y. Liu, J. Cheng, and J. Zhuo, "Fast stochastic variance reduced gradient method with momentum acceleration for machine learning," *arXiv:1703.07948*, 2017.

[35] L. Zhang, M. Mahdavi, and R. Jin, "Linear convergence with condition number independent access of full gradients," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 980–988.

[36] J. Konečný, J. Liu, P. Richtárik, , and M. Takáč, "Mini-batch semi-stochastic gradient descent in the proximal setting," *IEEE J. Sel. Top. Sign. Proces.*, vol. 10, no. 2, pp. 242–255, 2016.

[37] A. Nitanda, "Stochastic proximal gradient descent with acceleration techniques," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 1574–1582.

[38] G. Lan and Y. Zhou, "An optimal randomized incremental gradient method," *arXiv:1507.02000v3*, 2015.

[39] R. Frostig, R. Ge, S. M. Kakade, and A. Sidford, "Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 2540–2548.

[40] H. Lin, J. Mairal, and Z. Harchaoui, "A universal catalyst for first-order optimization," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 3366–3374.

[41] R. Babanezhad, M. O. Ahmed, A. Virani, M. Schmidt, J. Konecny, and S. Sallinen, "Stop wasting my gradients: Practical SVRG," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 2242–2250.

[42] Z. Allen-Zhu and Y. Yuan, "Improved SVRG for non-strongly-convex or sum-of-non-convex objectives," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 1080–1089.

[43] M. Frank and P. Wolfe, "An algorithm for quadratic programming," *Naval Res. Logist. Quart.*, vol. 3, pp. 95–110, 1956.

[44] E. Hazan and H. Luo, "Variance-reduced and projection-free stochastic optimization," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 1263–1271.

[45] F. Shang, Y. Liu, J. Cheng, K. W. Ng, and Y. Yoshida, "Variance reduced stochastic gradient descent with sufficient decrease," *arXiv:1703.06807*, 2017.

[46] L. Hien, C. Lu, H. Xu, and J. Feng, "Accelerated stochastic mirror descent algorithms for composite non-strongly convex optimization," *arXiv:1605.06892v2*, 2016.

[47] B. Woodworth and N. Srebro, "Tight complexity bounds for optimizing composite objectives," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 3639–3647.

[48] A. Defazio, "A simple practical accelerated method for finite sums," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 676–684.

[49] X. Li, T. Zhao, R. Arora, H. Liu, and J. Haupt, "Nonconvex sparse learning via stochastic optimization with progressive variance reduction," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 917–925.

[50] L. Nguyen, J. Liu, K. Scheinberg, and M. Takáč, "SARAH: A novel method for machine learning problems using stochastic recursive gradient," *arXiv:1703.00102v1*, 2017.

[51] L. W. Zhong and J. T. Kwok, "Fast stochastic alternating direction method of multipliers," in *Proc. 31st Int. Conf. Mach. Learn.*, 2014, pp. 46–54.

[52] S. Zheng and J. T. Kwok, "Fast-and-light stochastic ADMM," in *Proc. 25th Int. Joint Conf. Artif. Intell.*, 2016, pp. 2407–2613.

[53] S. J. Reddi, S. Sra, B. Poczos, and A. Smola, "Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 1145–1153.

[54] S. Reddi, A. Hefny, S. Sra, B. Poczos, and A. Smola, "On variance reduction in stochastic gradient descent and its asynchronous variants," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 2629–2637.

[55] J. D. Lee, Q. Lin, T. Ma, and T. Yang, "Distributed stochastic variance reduced gradient methods and a lower bound for communication complexity," *arXiv:1507.07595v2*, 2016.

[56] Y. Bengio, "Learning deep architectures for AI," *Found. Trends Mach. Learn.*, vol. 2, no. 1, pp. 1–127, 2009.

[57] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points  online stochastic gradient for tensor decomposition," in *Proc. 28th Conf. Learn. Theory*, 2015, pp. 797–842.

[58] A. Neelakantan, L. Vilnis, Q. V. Le, I. Sutskever, L. Kaiser, K. Kurach, and J. Martens, "Adding gradient noise improves learning for very deep networks," *arXiv:1511.06807*, 2015.

[59] N. Flammarion and F. Bach, "From averaging to acceleration, there is only a step-size," in *Proc. 28th Conf. Learn. Theory*, 2015, pp. 658–695.

[60] C. Tan, S. Ma, Y. Dai, and Y. Qian, "Barzilai-Borwein step size for stochastic gradient descent," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 685–693.

[61] Z. Allen-Zhu and E. Hazan, "Optimal black-box reductions between optimization objectives," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 1606–1614.

[62] B. Carpenter, "Lazy sparse stochastic gradient descent for regularized multinomial logistic regression," *Tech. Rep.*, 2008.

[63] J. Langford, L. Li, and T. Zhang, "Sparse online learning via truncated gradient," *J. Mach. Learn. Res.*, vol. 10, pp. 777–801, 2009.

[64] Y. Nesterov, "Smooth minimization of non-smooth functions," *Math. Program.*, vol. 103, pp. 127–152, 2005.

[65] Y. Xu, Y. Yan, Q. Lin, and T. Yang, "Homotopy smoothing for non-smooth problems with lower complexity than $O(1/\epsilon)$," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 1208–1216.

[66] P. Zhao and T. Zhang, "Stochastic optimization with importance sampling for regularized loss minimization," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 1–9.

[67] D. Needell, N. Srebro, and R. Ward, "Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm," *Math. Program.*, vol. 155, pp. 549–573, 2016.

[68] J. Konečný and P. Richtárik, "Semi-stochastic gradient descent methods," *ArXiv Preprint: 1312.1666v2*, 2015.

[69] G. Lan, "An optimal method for stochastic composite optimization," *Math. Program.*, vol. 133, pp. 365–397, 2012.

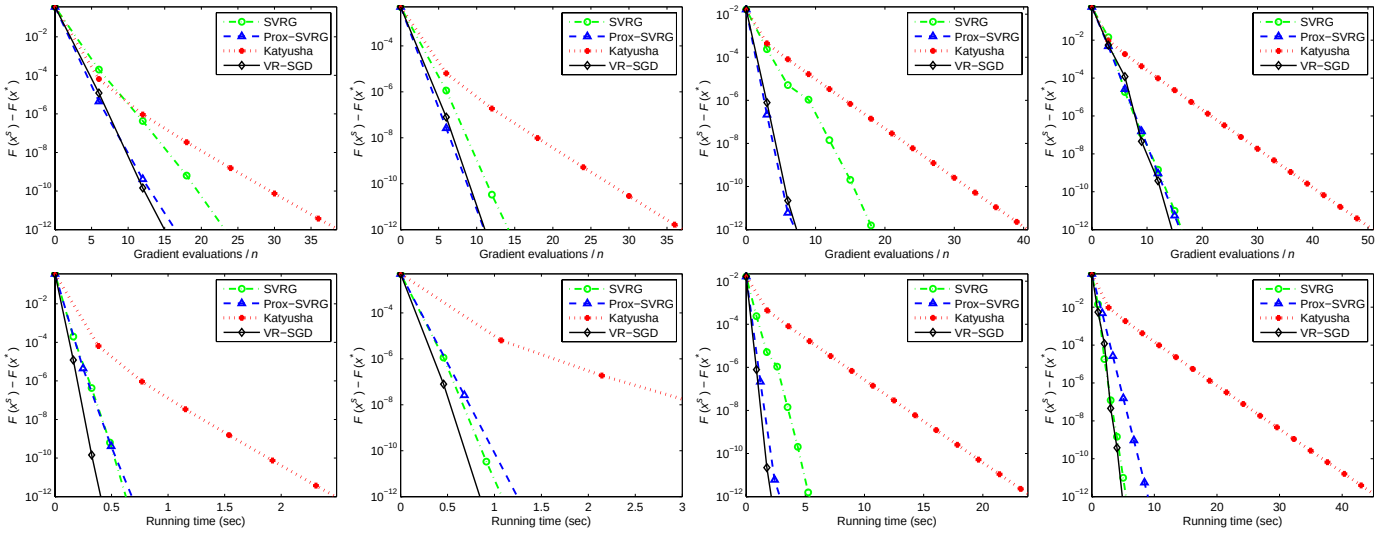


(a) Adult: $\lambda_1 = 10^{-4}$ (b) Protein: $\lambda_1 = 10^{-4}$ (c) Covtype: $\lambda_1 = 10^{-4}$ (d) Sido0: $\lambda_1 = 5 * 10^{-3}$

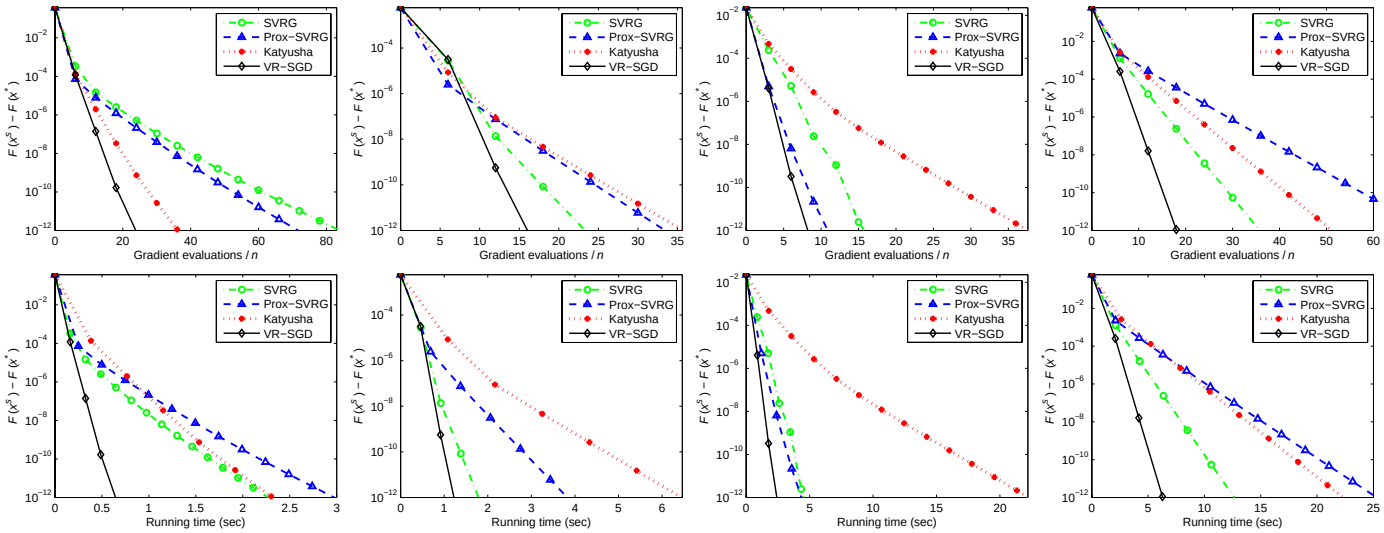


(e) Adult: $\lambda_1 = 10^{-5}$ (f) Protein: $\lambda_1 = 10^{-5}$ (g) Covtype: $\lambda_1 = 10^{-5}$ (h) Sido0: $\lambda_1 = 10^{-4}$

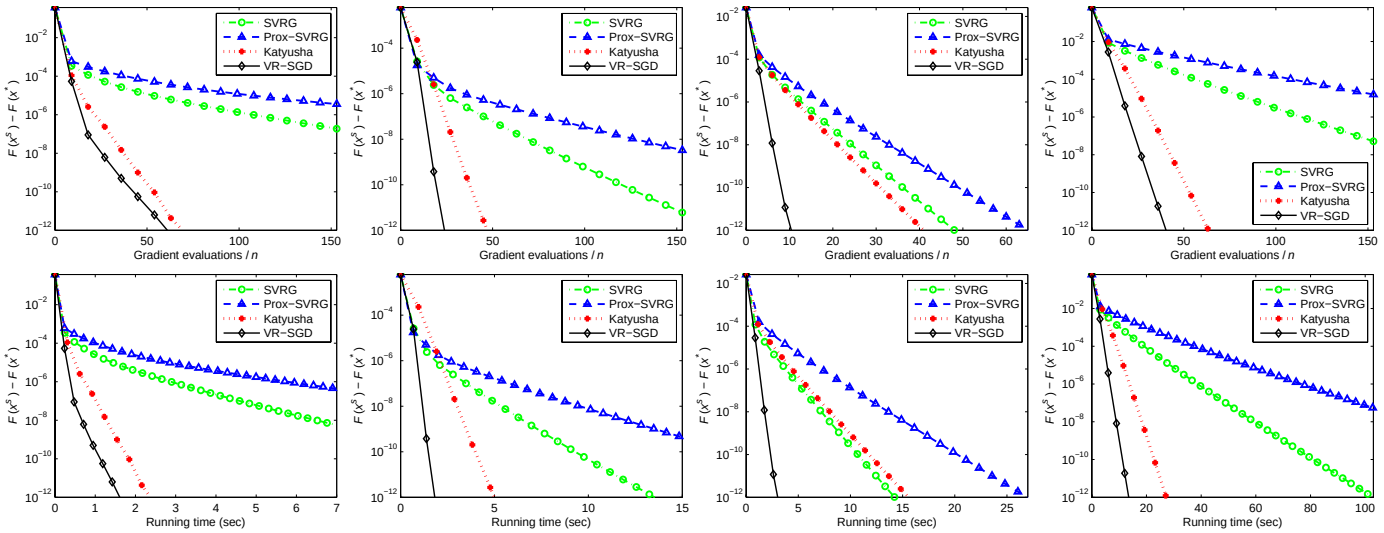


(i) Adult: $\lambda_1 = 10^{-6}$ (j) Protein: $\lambda_1 = 10^{-6}$ (k) Covtype: $\lambda_1 = 10^{-6}$ (l) Sido0: $\lambda_1 = 10^{-5}$

Fig. 5. Comparison of SVRG [1], Prox-SVRG [2], Katyusha [3], and VR-SGD for solving $\ell_2$-norm regularized logistic regression problems (i.e., $\lambda_2 = 0$). In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes (top) or running time (bottom).

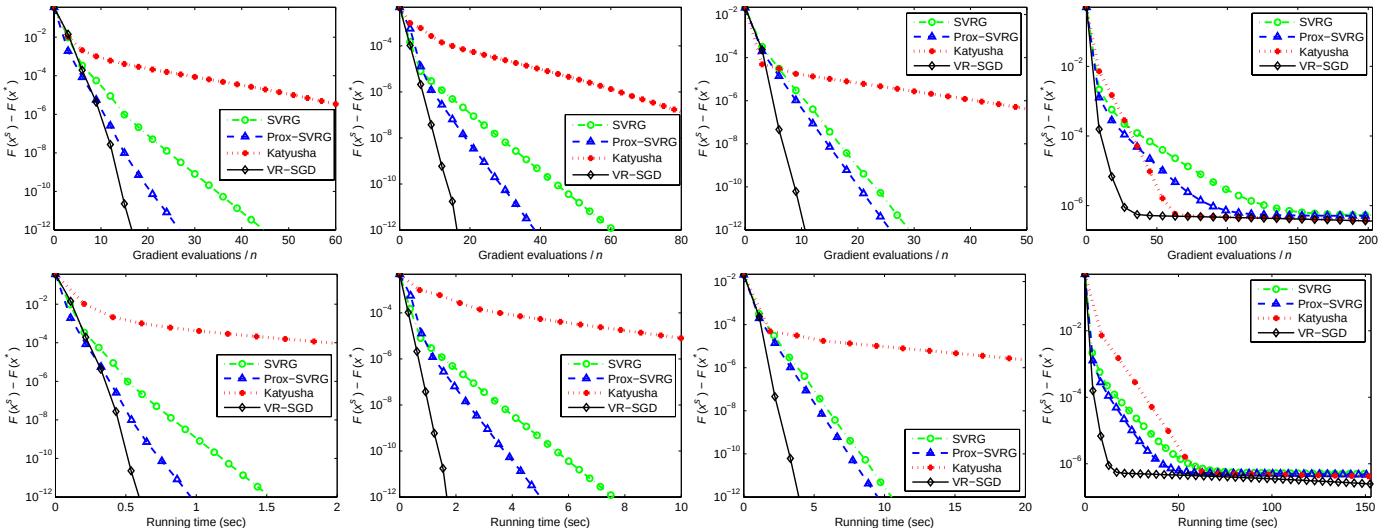


(a) $\lambda_2 = 10^{-4}$ (b) $\lambda_2 = 10^{-3}$

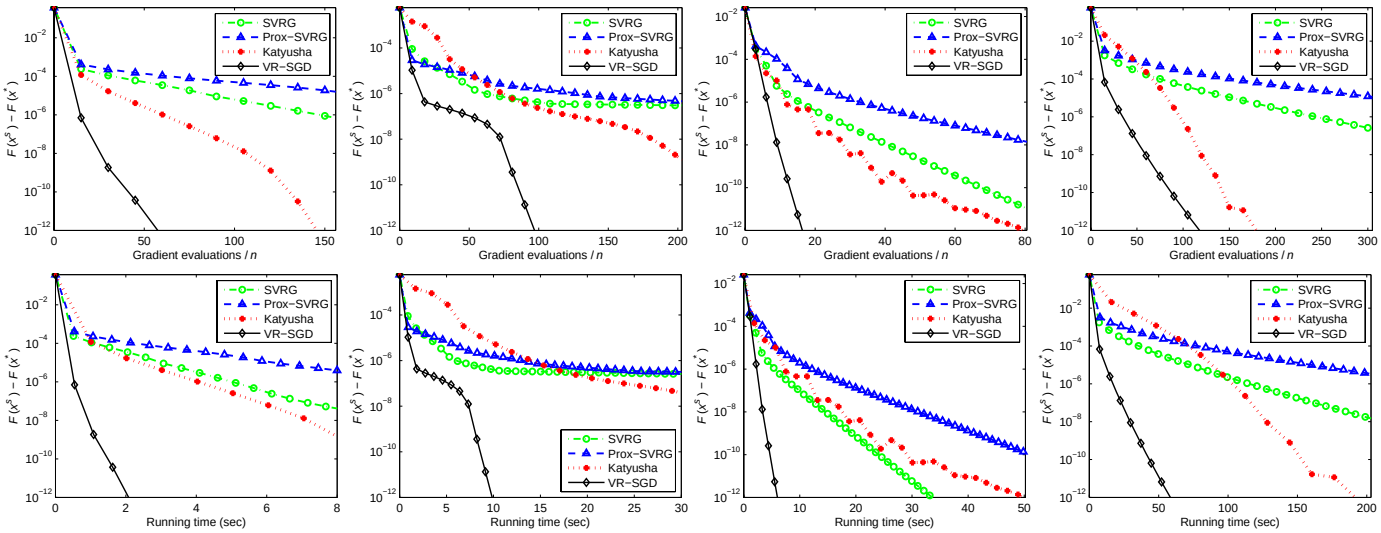


(c) $\lambda_2 = 10^{-5}$ (d) $\lambda_2 = 10^{-4}$

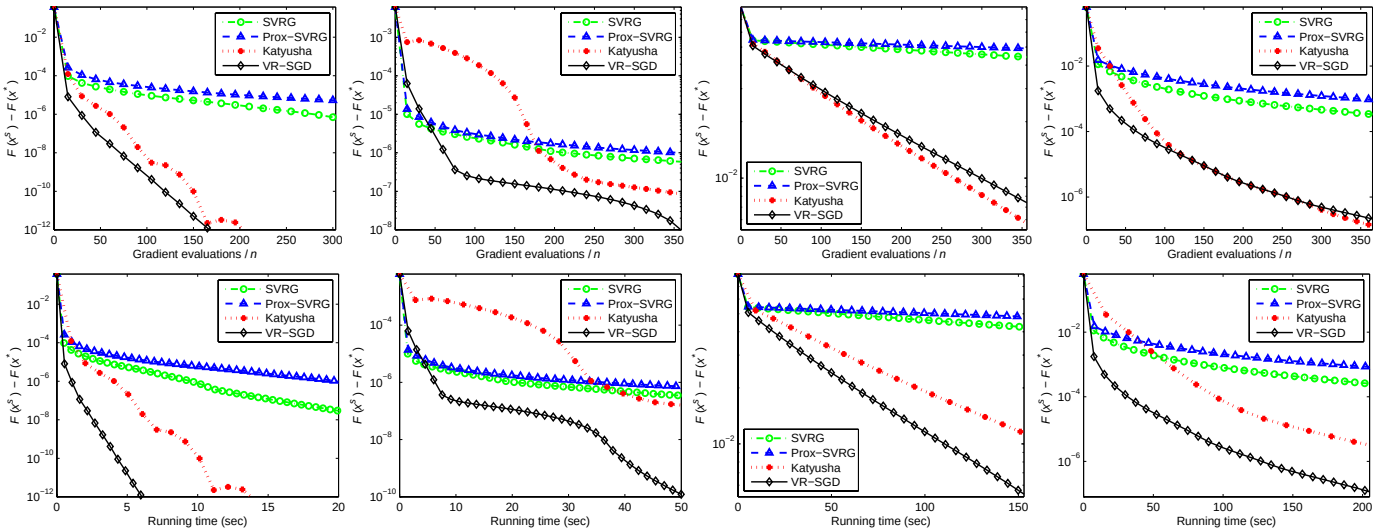


(e) $\lambda_2 = 10^{-6}$ (f) $\lambda_2 = 10^{-5}$

Fig. 6. Comparison of SVRG, Prox-SVRG [2], Katyusha [3], and VR-SGD for $\ell_1$-norm regularized logistic regression problems (i.e., $\lambda_1 = 0$) on the four data sets: Adult (the first column), Protein (the sconced column), Covtype (the third column), and Sido0 (the last column). In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes (top) or running time (bottom).

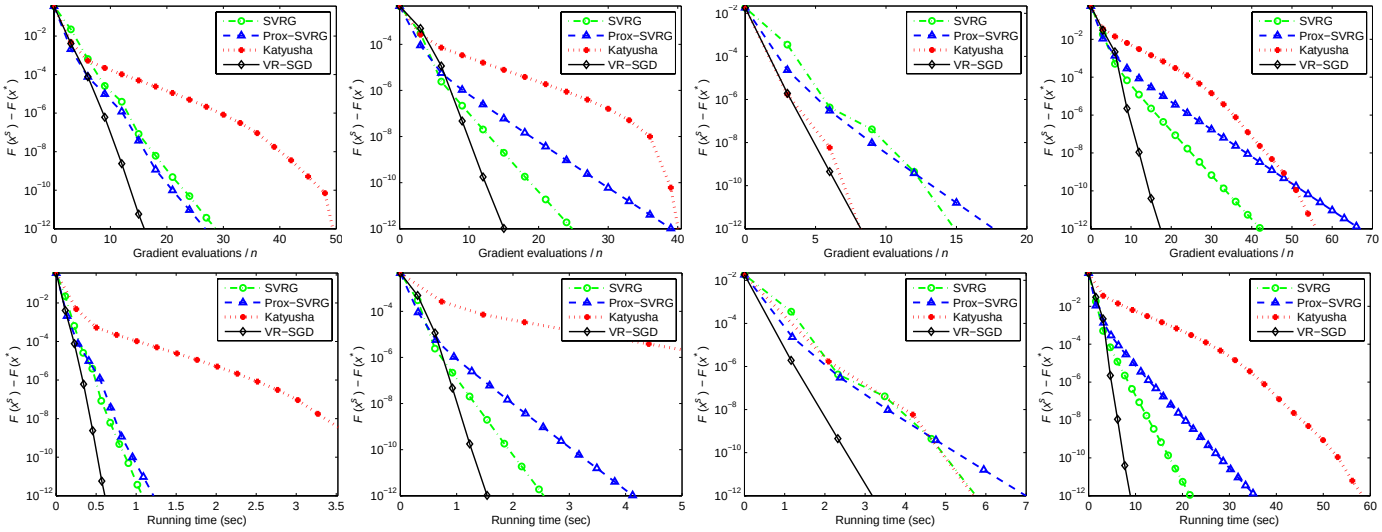


(a) $\lambda_1 = 10^{-5}$ and $\lambda_2 = 10^{-4}$ (b) $\lambda_1 = 10^{-4}$ and $\lambda_2 = 10^{-5}$

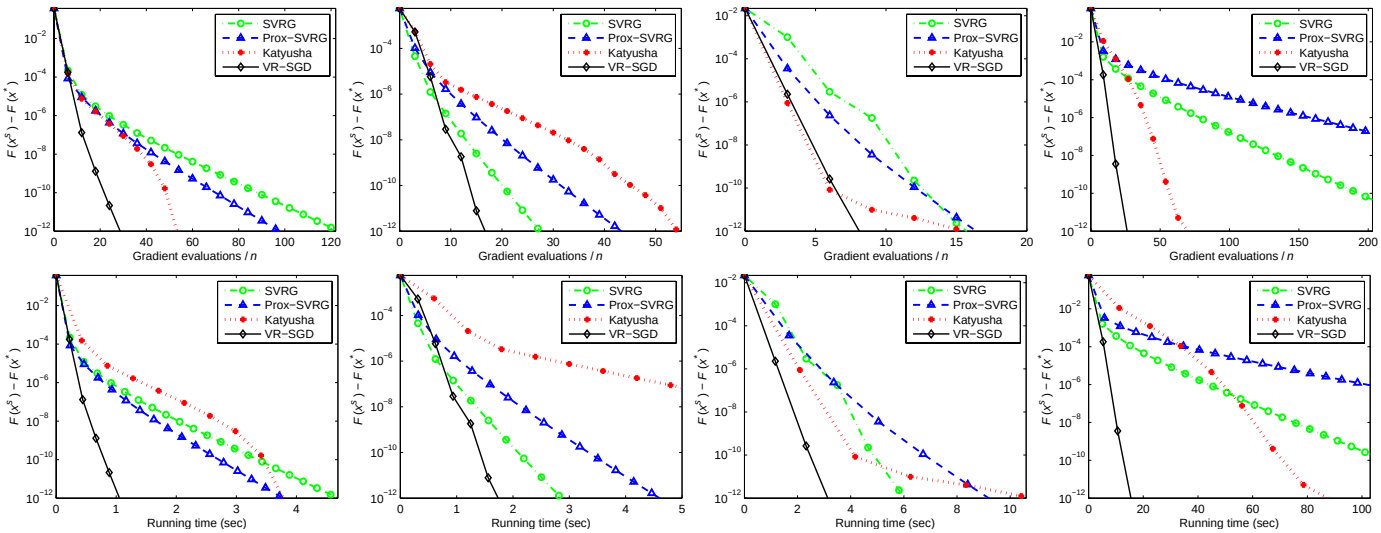


(c) $\lambda_1 = 10^{-5}$ and $\lambda_2 = 10^{-5}$ (d) $\lambda_1 = 10^{-5}$ and $\lambda_2 = 10^{-4}$

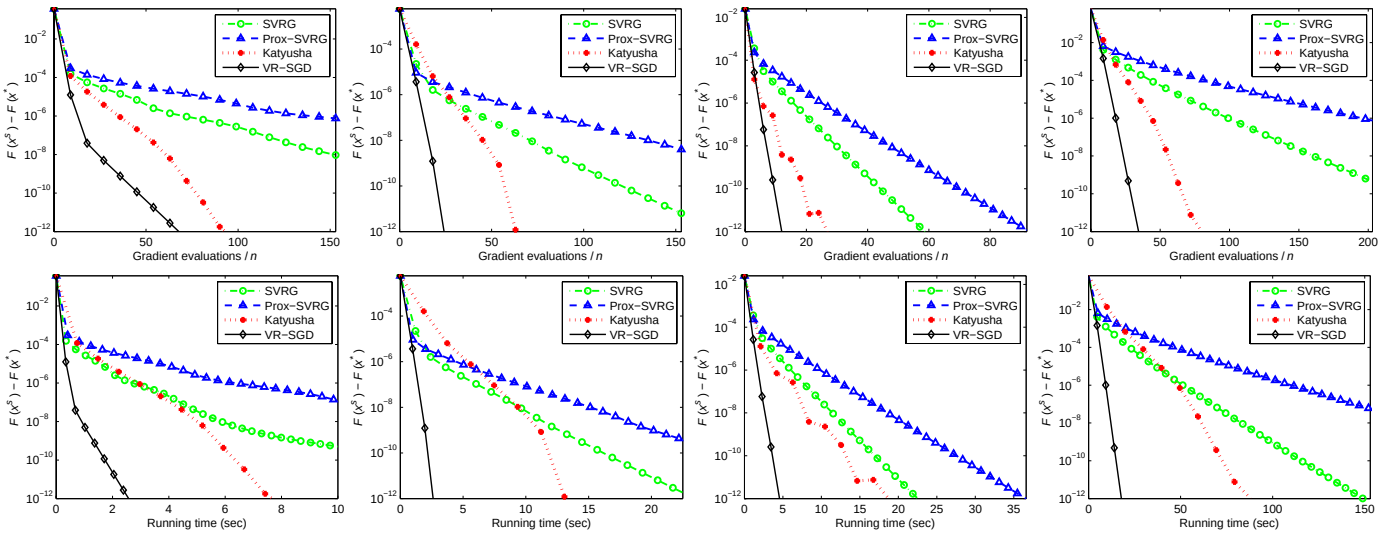


(e) $\lambda_1 = 10^{-6}$ and $\lambda_2 = 10^{-5}$ (f) $\lambda_1 = 10^{-5}$ and $\lambda_2 = 10^{-5}$

Fig. 7. Comparison of SVRG, Prox-SVRG [2], Katyusha [3], and VR-SGD for solving elastic net regularized logistic regression problems on the four data sets: Adult (the first column), Protein (the sconced column), Covtype (the third column), and Sido0 (the last column). In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes (top) or running time (bottom).

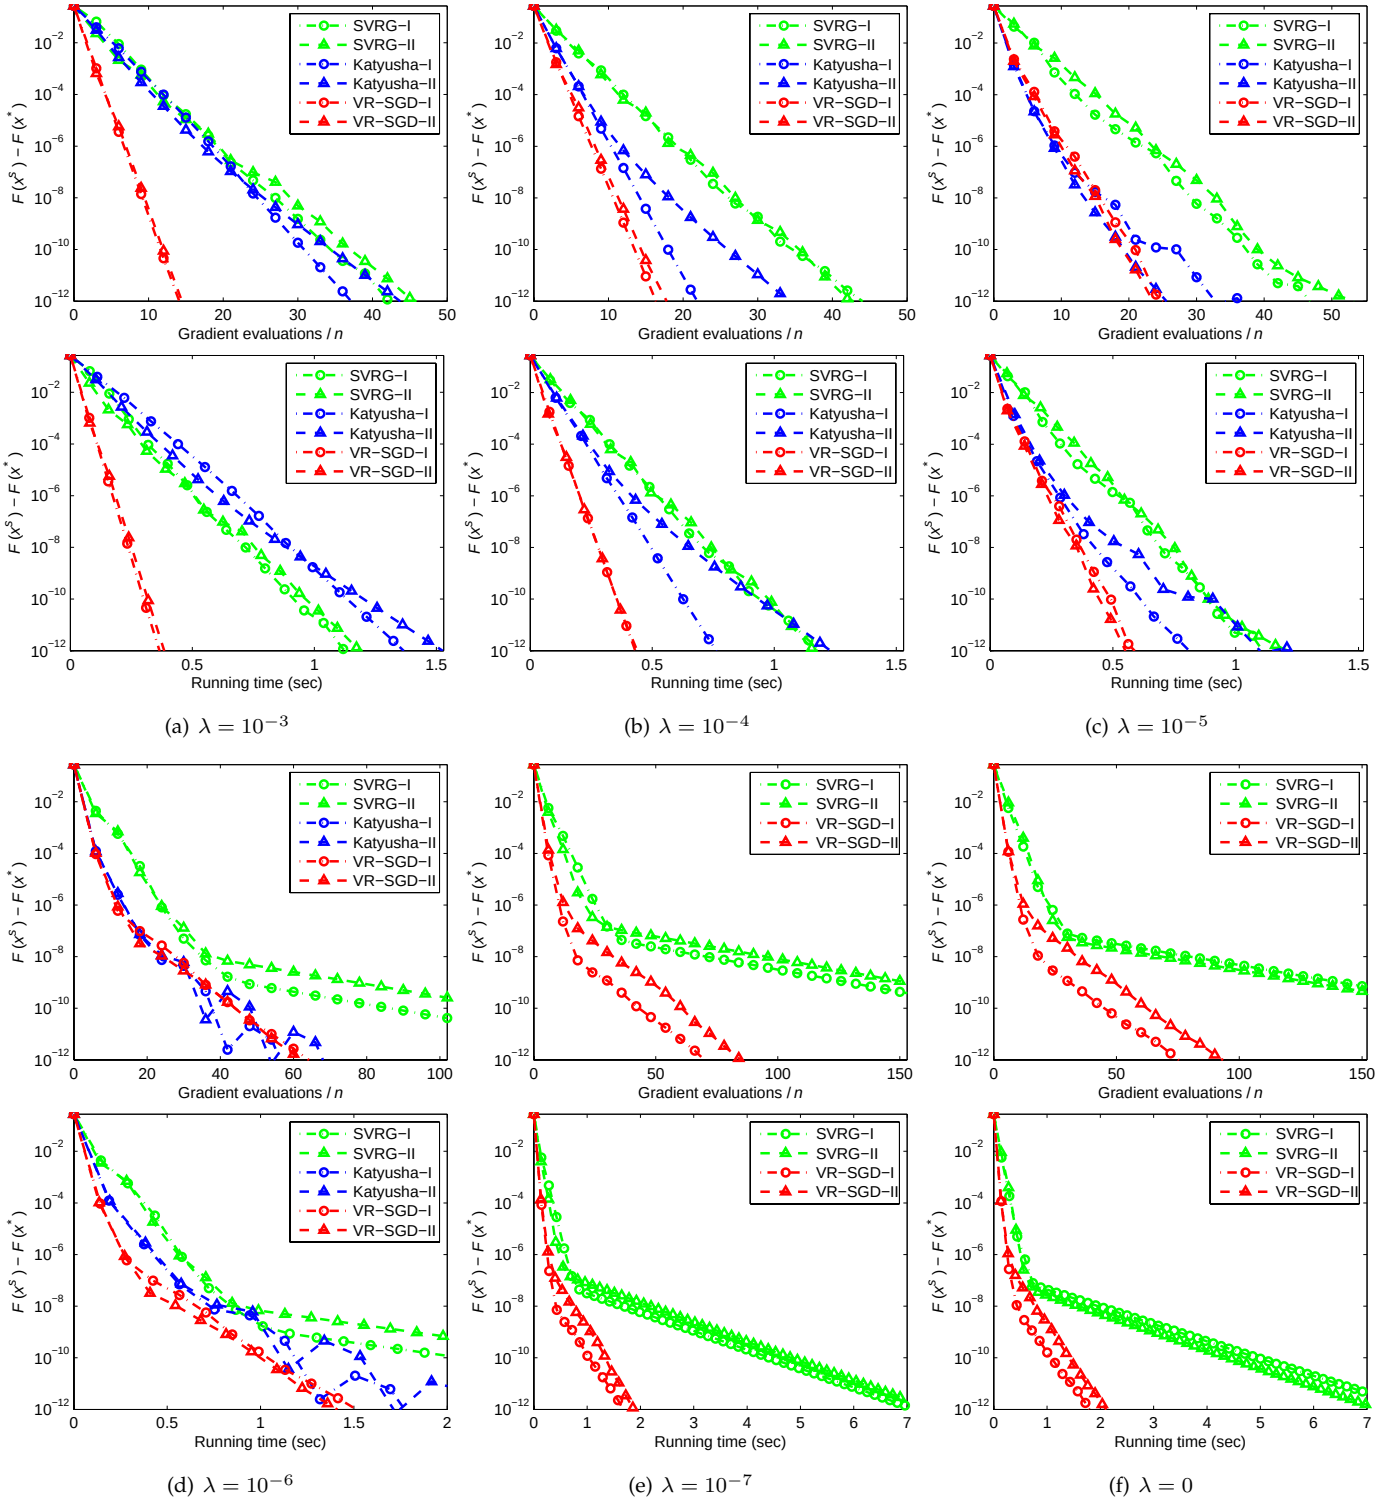


(a) $\lambda = 10^{-3}$ (b) $\lambda = 10^{-4}$ (c) $\lambda = 10^{-5}$

(d) $\lambda = 10^{-6}$ (e) $\lambda = 10^{-7}$ (f) $\lambda = 0$

Fig. 8. Comparison of SVRG [1], Katyusha [3], VR-SGD and their proximal versions for solving ridge regression problems with different regularization parameters on the Adult data set. In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes (top) or running time (bottom).

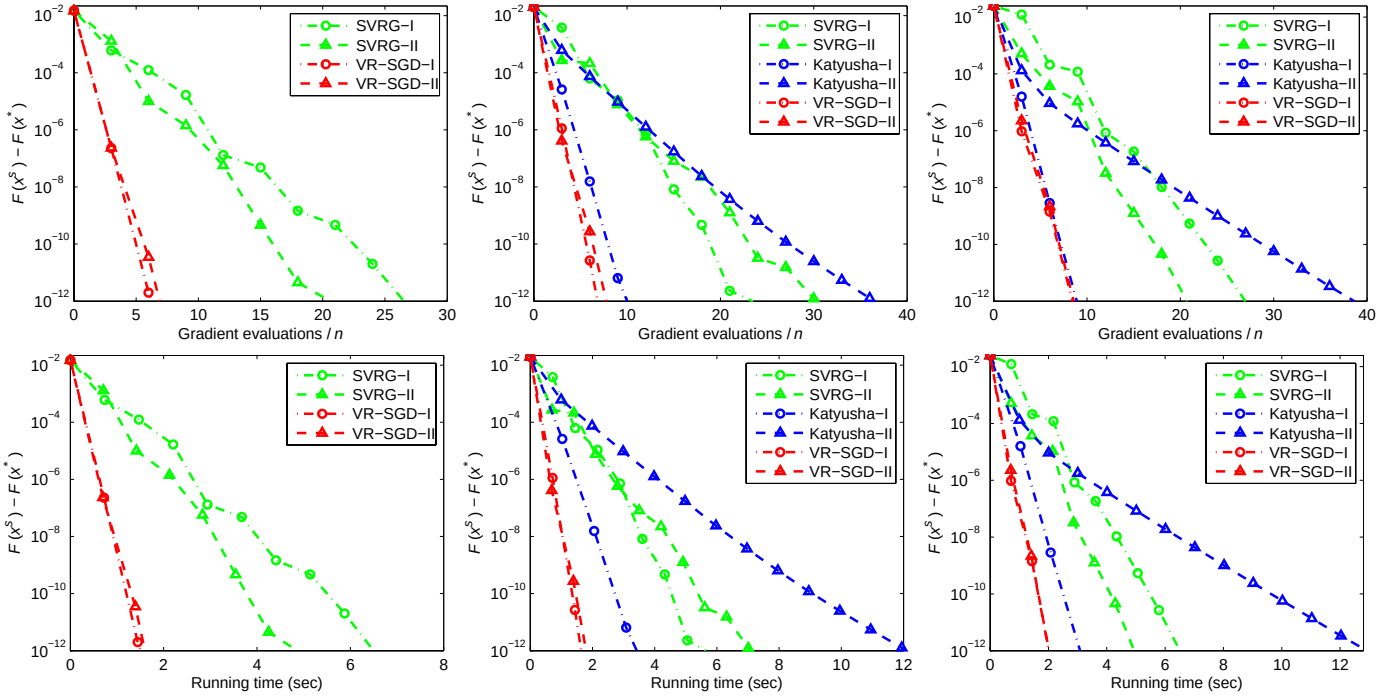


(a) $\lambda = 10^{-3}$ (b) $\lambda = 10^{-4}$ (c) $\lambda = 10^{-5}$

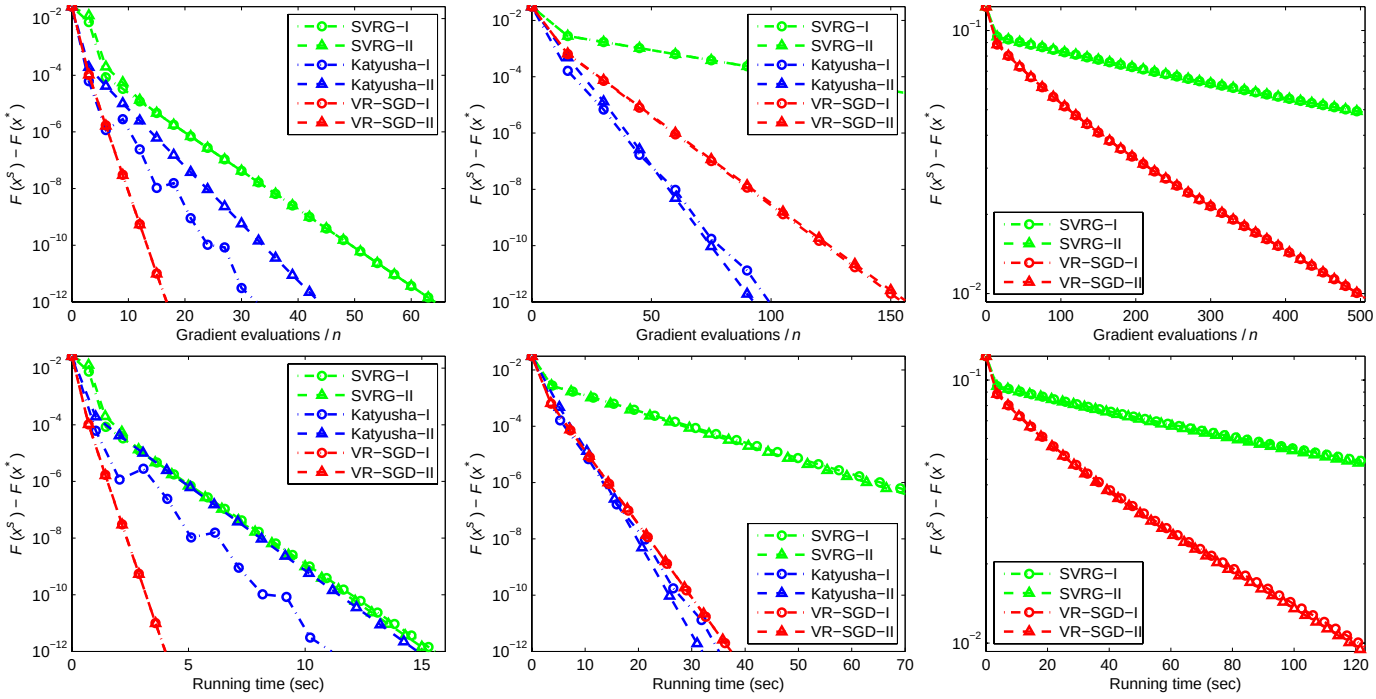


(d) $\lambda = 10^{-6}$ (e) $\lambda = 10^{-7}$ (f) $\lambda = 0$

Fig. 9. Comparison of SVRG [1], Katyusha [3], VR-SGD and their proximal versions for solving ridge regression problems with different regularization parameters on the Covtype data set. In each plot, the vertical axis shows the objective value minus the minimum, and the horizontal axis is the number of effective passes (top) or running time (bottom).

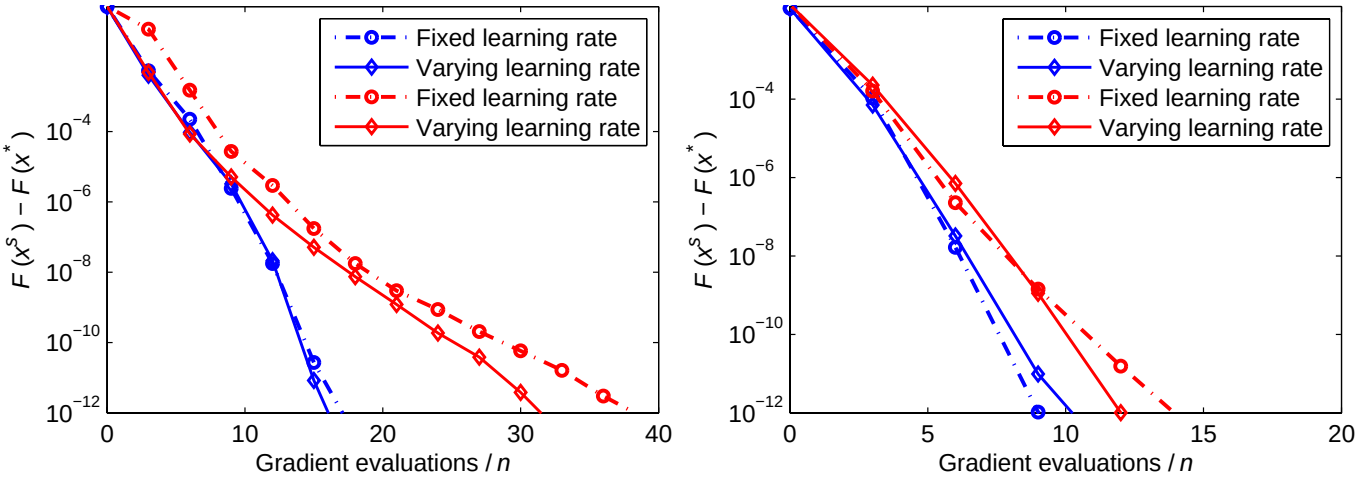


(a) $\ell_1$-norm regularized logistic regression: Adult (left) and Covtype (right)

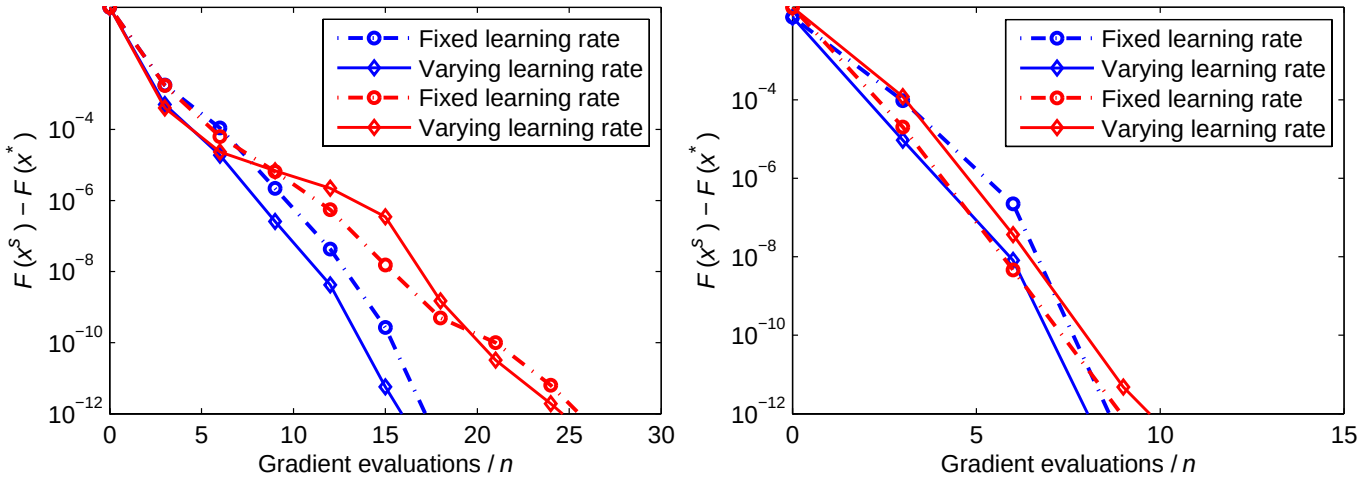


(b) Lasso: Adult (left) and Covtype (right)

Fig. 10. Comparison of Algorithm 3 with fixed and varying learning rates for solving $\ell_1$-norm (i.e., $\lambda\|x\|_1$) regularized logistic regression and Lasso problems with $\lambda = 10^{-4}$ (blue lines) and $\lambda = 10^{-5}$ (red lines, best viewed in colors). Note that the regularization parameter is set to $10^{-3}$ and $10^{-4}$ for solving Lasso problems on the Covtype data set.